arXiv:1412.5655v1 [cs.CC] 17 Dec 2014

New algorithms and lower bounds for monotonicity testing

Xi Chen*
Columbia University

Rocco A. Servedio†
Columbia University

Li-Yang Tan‡
Simons Institute, UC Berkeley

{xichen,rocco,liyang}@cs.columbia.edu

December 19, 2014

Abstract

We consider the problem of testing whether an unknown Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ is monotone versus ε-far from every monotone function. The two main results of this paper are a new lower bound and a new algorithm for this well-studied problem.

Lower bound: We prove an $\tilde{\Omega}(n^{1/5})$ lower bound on the query complexity of any non-adaptive two-sided error algorithm for testing whether an unknown Boolean function f is monotone versus constant-far from monotone. This gives an exponential improvement on the previous lower bound of $\Omega(\log n)$ due to Fischer *et al.* [FLN+02]. We show that the same lower bound holds for monotonicity testing of Boolean-valued functions over hypergrid domains $\{1,\ldots,m\}^n$ for all $m \geq 2$.

Upper bound: We give an $\tilde{O}(n^{5/6})\mathrm{poly}(1/\varepsilon)$-query algorithm that tests whether an unknown Boolean function f is monotone versus ε-far from monotone. Our algorithm, which is non-adaptive and makes one-sided error, is a modified version of the algorithm of Chakrabarty and Seshadhri [CS13a], which makes $\tilde{O}(n^{7/8})\mathrm{poly}(1/\varepsilon)$ queries.

*Supported by NSF grant CCF-1149257 and a Sloan research fellowship.

†Supported by NSF grants CCF-1115703 and CCF-1319788.

‡This work was done while the author was at Columbia Univerisity, supported by NSF grants CCF-1115703 and CCF-1319788.

1 Introduction

Monotonicity is a basic and natural property of functions. In the field of property testing, the problem of efficiently testing whether an unknown function is monotone has been the focus of a long and fruitful line of research, with many works (see e.g. [GGLR98, DGL+99, GGL+00, EKK+00, FLN+02, Fis04, BKR04, ACCL07, HK08, RS09, BBM12, BCGSM12, RRS+12, CS13a, CS13b, CS13c, BRY13]) studying this problem for functions with various domains and ranges.

In this work we will be concerned with the classical problem of testing monotonicity of *Boolean functions* $f : \{-1,1\}^n \to \{-1,1\}$, which was first posed and considered explicitly by Goldreich *et al.* [GGLR98]. Recall that a Boolean function f is monotone if $f(x) \leq f(y)$ for all $x \prec y$, where $\prec$ denotes the bitwise partial order on the hypercube. Let $\mathrm{dist}(f,g) := \mathbf{Pr}_{\boldsymbol{x}\in\{-1,1\}^n}[f(\boldsymbol{x}) \neq g(\boldsymbol{x})]$; we say that f is *ε-close to monotone* if $\mathrm{dist}(f,g) \leq \varepsilon$ for some monotone Boolean function g, and that f is *ε-far from monotone* otherwise. We will be interested in query-efficient randomized testing algorithms for the following task:

> *Given as input a distance parameter $\varepsilon > 0$ and oracle access to an unknown Boolean function $f : \{-1,1\}^n \to \{-1,1\}$, output* Yes *with probability at least $2/3$ if f is monotone, and* No *with probability at least $2/3$ if f is ε-far from monotone.*

The work of Goldreich *et al.* [GGLR98] proposed a simple "edge tester" which queries uniform random edges of $\{-1,1\}^n$ hoping to find an edge whose endpoints violate monotonicity. [GGLR98] proved an $O(n^2 \log(1/\varepsilon)/\varepsilon)$ upper bound on the query complexity of the edge tester, which was subsequently improved to $O(n/\varepsilon)$ in the journal version [GGL+00]. Fischer *et al.* [FLN+02] established the first lower bounds shortly after, showing that there exists a constant distance parameter $\varepsilon_0 > 0$ such that $\Omega(\log n)$ queries are necessary for any *non-adaptive* tester (one whose queries do not depend on the oracle's responses to prior queries). This directly implies an $\Omega(\log\log n)$ lower bound for adaptive testers, since any q-query adaptive tester can be simulated by a non-adaptive one that simply carries out all 2^q possible executions. These upper and lower bounds were the best known for more than a decade, until the recent work of Chakrabarty and Seshadhri [CS13a] improved on the linear upper bound of Goldreich *et al.* with an $\tilde{O}(n^{7/8}\varepsilon^{-3/2})$-query tester.

Our main contributions in this work are (i) a new lower bound that improves on the [FLN+02] lower bound by an exponential factor, and (ii) a new algorithm that improves on the [CS13a] upper bound (in terms of the dependence on n) by a polynomial factor. We now describe these contributions in more detail.

Our lower bound. We give an exponential improvement on the above-mentioned lower bounds of Fischer *et al.*:

Theorem 1. *There exists a universal constant $\varepsilon_0 > 0$ such that any non-adaptive algorithm for testing whether an unknown Boolean function is monotone versus ε_0-far from monotone must make $\Omega(n^{1/5}(\log n)^{-2/5})$ queries. Consequently, any adaptive algorithm must make $\Omega(\log n)$ queries.*

While the aforementioned results of Fischer *et al.* represent the previous best lower bounds on the general testing problem as defined above, additional lower bounds are known for several restricted versions of the problem. In the same paper Fischer *et al.* gave an $\Omega(\sqrt{n})$ lower bound on the query complexity of any non-adaptive *one-sided* tester, i.e. one that always outputs Yes when f is monotone (again, this directly implies an $\Omega(\log n)$ lower bound for adaptive one-sided testers). Restricting further, a *pair tester* is a non-adaptive one-sided tester that independently draws pairs

of comparable points $x \prec y$ from some distribution and rejects if and only if some pair that is drawn violates monotonicity. Briët *et al.* [BCGSM12] proved an $\Omega(n/(\varepsilon \log n))$ lower bound on the query complexity of pair testers whose query complexity can be written as $q(n)/\varepsilon$ for some function q.

In addition to Theorem 1, we show that essentially the same lower bound holds for monotonicity testing of Boolean-valued functions over hypergrid domains $\{1, \dots, m\}^n$ for $m \geq 2$. (Below and throughout this paper we write $[m]$ to denote $\{1, 2, \dots, m\}$.) Our most general lower bound is the following:

Theorem 2. *There exists a universal constant $\varepsilon_0 > 0$ such that for all $m \geq 2$, any non-adaptive algorithm for testing whether an unknown function $f : [m]^n \to \{-1, 1\}$ is monotone versus ε_0-far from monotone must make $\tilde{\Omega}(n^{1/5})$ queries.*

To the best of our knowledge Theorem 2 is the first lower bound for testing monotonicity of *Boolean-valued* functions over hypergrid domains. Recent papers of Chakrabarty and Seshadhri [CS13b, CS13c] and Blais *et al.* [BRY13] essentially close the problem of testing monotonicity of functions $f : [m]^n \to \mathbb{N}$, showing that $\Theta(n \log m)$ queries are both necessary and sufficient; however, their lower bounds crucially depend on the functions considered having range $\mathbb{N}$ rather than $\{-1, 1\}$.

Our algorithm. We present a new algorithm for monotonicity testing and prove the following result about its performance:

Theorem 3. *There is a $\tilde{O}(n^{5/6}\varepsilon^{-4})$-query one-sided non-adaptive algorithm for testing whether an unknown n-variable Boolean function is monotone versus ε-far from monotone.*

Recall that the one-sided, non-adaptive tester of Chakrabarty and Seshadhri [CS13a] makes $\tilde{O}(n^{7/8}\varepsilon^{-3/2})$ queries. Thus, while the query complexity of our tester is worse as a function of $1/\varepsilon$ (though still polynomial), its query complexity is polynomially better as a function of n.[1] Like the [CS13a] algorithm, our algorithm is a pair tester, but it evades the $\Omega(n/(\varepsilon \log n))$ lower bound of [BCGSM12] because its query complexity is not of the form $q(n)/\varepsilon$. Our algorithm builds on the tools developed in [CS13a]; its high-level structure is similar to that of the [CS13a] algorithm, but with an important difference that enables an improved analysis. See Section 1.2 for more discussion on this point.

1.1 The lower bound approach

Our lower bound for testing monotonicity builds on previous lower bounds for testing restricted classes of *linear threshold functions* (LTFs). Recall that $f : \{-1, 1\}^n \to \{-1, 1\}$ is a linear threshold function if there exist $w_1, \dots, w_n, \theta \in \mathbb{R}^n$ such that $f(x) = \text{sign}(w \cdot x - \theta)$ for all $x \in \{-1, 1\}^n$.

Background. A *signed majority function* is a linear threshold function of the special form $f(x) = \text{sign}(w \cdot x)$ where $w \in \{-1, 1\}^n$. While [MORS10] showed that the class of all LTFs is ε-testable using $\text{poly}(1/\varepsilon)$ queries (independent of n), in [MORS09] Matulef *et al.* gave an $\Omega(\log n)$ lower bound for non-adaptive algorithms that ε_0-test whether $f : \{-1, 1\}^n \to \{-1, 1\}$ is a signed majority function, where $\varepsilon_0 > 0$ is a universal constant. Like many lower bound arguments in

[1] Recall that in property testing the dependence on the size parameter "n" is typically viewed as more important than the dependence on the "closeness" parameter ε. Indeed, ε is often viewed as a constant, so testers with query complexities that are exponential (or worse) as a function of $1/\varepsilon$ but independent of n are commonly referred to as "constant-query testers."

property testing, the proof of [MORS09] employs Yao's minimax principle [Yao77], and works by exhibiting two distributions $\mathcal{D}_{yes}$ and $\mathcal{D}_{no}$ over LTFs — more precisely, $\mathcal{D}_{yes}$ is the uniform distribution over all 2^n signed majority functions, and $\mathcal{D}_{no}$ is the uniform distribution over a set of LTFs almost all of which are constant-far from every signed majority function — and arguing that for $q = o(\log n)$, any deterministic q-query algorithm cannot distinguish between the two distributions with non-negligible success probability. (We note that a typical function from $\mathcal{D}_{yes}$ is far from being monotone, and that the same holds for a typical LTF drawn from the $\mathcal{D}_{no}$ distribution of [MORS09].) A key tool in the [MORS09] proof is the Berry–Esséen "central limit theorem (CLT) with error bounds" for sums of independent real-valued random variables.

An *embedded majority function of size* k is an LTF $f : \{-1,1\}^n \to \{-1,1\}$ of the form $f(x) = \mathrm{sign}(w \cdot x)$ where $w \in \{0,1\}^n$ is a vector with exactly k ones. In [BO10] Blais and O'Donnell showed that for $k = n/2$, any non-adaptive testing algorithm for the class of all embedded majority functions of size exactly $n/2$ must make $\Omega(n^{1/12})$ queries. Their proof employed a $\mathcal{D}_{yes}$ distribution which is the uniform distribution over all embedded majority functions of size $n/2$, and a $\mathcal{D}_{no}$ distribution which is supported on certain monotone LTFs (which are far from embedded majority functions of size $n/2$). A key technical ingredient in the proofs of [BO10] is a multidimensional extension of the Berry–Esséen theorem (to independent sums of $\mathbb{R}^q$-valued random variables) which was essentially established in the work of [GOWZ10], building on ingredients from [Mos08]. Subsequently Ron and Servedio [RS13] adapted the arguments of [BO10] to give an improved analysis of the same $\mathcal{D}_{yes}$ and $\mathcal{D}_{no}$ distributions from [MORS09] and establish an $\Omega(n^{1/12})$-query lower bound for non-adaptive algorithms that ε_0-test whether $f : \{-1,1\}^n \to \{-1,1\}$ is a signed majority function, thus exponentially improving over the [MORS09] lower bounds for this problem.

This work. Neither the [BO10] construction nor the [MORS09, RS13] construction can be used directly to establish a lower bound for monotonicity testing of functions $f : \{-1,1\}^n \to \{-1,1\}$; as described above, in the [BO10] construction both the $\mathcal{D}_{yes}$ and $\mathcal{D}_{no}$ functions are monotone, and in the [MORS09, RS13] construction a typical function from either distribution is far from monotone. Nevertheless, in this work we show that ingredients from [BO10, RS13] can be leveraged to obtain a polynomial lower bound for testing monotonicity of functions $f : \{-1,1\}^n \to \{-1,1\}$. Like these earlier works we employ Yao's principle: we define a $\mathcal{D}_{yes}$ distribution that is supported on monotone LTFs, and a $\mathcal{D}_{no}$ distribution over LTFs that is almost entirely supported on LTFs that are constant-far from every monotone function, and use an analysis which is fairly similar to that of [BO10, RS13], to prove Theorem 1. Using the multidimensional Berry–Esséen theorem of [GOWZ10] to analyze our $\mathcal{D}_{yes}$ and $\mathcal{D}_{no}$ distributions would result in an $\Omega(n^{1/12})$ lower bound. To obtain our improved $\Omega(n^{1/5} \log^{-2/5} n)$ lower bound, we instead adapt a multidimensional CLT of Valiant and Valiant [VV11] (for Wasserstein distance) to our context.

1.2 The approach of our algorithm

Our algorithm builds on ingredients from [CS13a], so to explain our approach we first recall the necessary ingredients from that work. Fix a Boolean function[2] $f : \{0,1\}^n \to \{0,1\}$, and let us say that a pair of inputs (x, y) with $x \prec y$ is a *violated edge* if $f(x) = 1, f(y) = 0$ and (x, y) is an edge in $\{0,1\}^n$ (i.e. the Hamming distance between them is 1). [CS13a] establishes a very useful "dichotomy theorem" about Boolean functions $f : \{0,1\}^n \to \{0,1\}$ that are ε-far from monotone:

[2] For our algorithmic result it will be more convenient to view Boolean functions as mapping $\{0,1\}^n$ to $\{0,1\}$.

for any $s > 0$, any such function either must have $\Omega(\varepsilon s 2^n)$ violated edges, or must have a *matching* (i.e. a vertex-disjoint set) of $\Omega(\varepsilon 2^n/s)$ violated edges.

To use this dichotomy theorem, Chakrabarty and Seshadhri [CS13a] define a "path tester" which works essentially as follows: it selects a random directed path $\boldsymbol{p}$ of n edges from 0^n up to 1^n, draws two uniform random points $\boldsymbol{x} \prec \boldsymbol{y}$ from the "middle layers" of $\boldsymbol{p}$, and rejects if $\boldsymbol{x}$ and $\boldsymbol{y}$ violate monotonicity, i.e. $f(\boldsymbol{x}) = 1$ and $f(\boldsymbol{y}) = 0$.[3] They prove that if f has a matching of $\Omega(\sigma 2^n)$ violated edges, then their path tester will uncover a violation and reject with probability $\tilde{\Omega}(\sigma^3/\sqrt{n})$. (Roughly speaking, they show that about an $\Omega(\sigma)$ fraction of possible outcomes of $\boldsymbol{y}$, corresponding to the $\sigma 2^n$ upper endpoints of the edges in the matching, are such that with probability $\tilde{\Omega}(\sigma^2/\sqrt{n})$ over the random draw of $\boldsymbol{x}$, the pair $\boldsymbol{y}$ and $\boldsymbol{x}$ together constitute a violation.) On the other hand, if f does not have a matching of this size then (by the dichotomy theorem) it must have $\Omega((\varepsilon^2/\sigma)2^n)$ violated edges, so the edge tester of [GGLR98] (querying the endpoints of a uniform random edge) will hit a violated edge with probability $\Omega(\varepsilon^2/(\sigma n))$. Their final algorithm runs their path tester with probability $1/2$ and queries a random edge with probability $1/2$. Choosing σ suitably to equalize the two rejection probabilities, this is a two-query algorithm which succeeds in uncovering a violation for any ε-far-from-monotone function f with probability $\tilde{\Omega}(\varepsilon^{3/2}/n^{7/8})$, giving them a one-sided non-adaptive tester which makes $\tilde{O}(n^{7/8}/\varepsilon^{3/2})$ queries overall.

Our algorithm follows the same high-level framework described above, but differs from [CS13a] by employing a different path tester. After selecting a random path $\boldsymbol{p}$, instead of (essentially) drawing two independent uniform points from the middle layers of the path as is done in [CS13a], our path tester draws a *correlated* pair of points from $\boldsymbol{p}$. More precisely, it selects the first point $\boldsymbol{y}$ independently from the middle layers of $\boldsymbol{p}$, and preferentially selects the second point $\boldsymbol{x}$ from $\boldsymbol{p}$ in a way which favors points which are closer to $\boldsymbol{y}$. Via a careful analysis we are able to show that if f has a matching of $\Omega(\sigma 2^n)$ violated edges, then our path tester will uncover a violation and reject with probability $\tilde{\Omega}(\sigma^2/\sqrt{n}) \cdot \mathrm{poly}(\varepsilon)$. Roughly speaking, we show that if $\boldsymbol{y}$ is a uniform random upper endpoint of the $\sigma 2^n$ edges in the matching (which occurs with probability about σ), then the probability that our tester selects a string $\boldsymbol{x}$ which gives a violation with $\boldsymbol{y}$ is $\tilde{\Omega}(\sigma/\sqrt{n}) \cdot \mathrm{poly}(\varepsilon)$. Trading this off against the success probability of the edge tester using the dichotomy theorem, we obtain our improved query bound.

Organization of this paper. Our lower bound results are established Sections 2 through 4. The two distributions $\mathcal{D}_{yes}$ and $\mathcal{D}_{no}$ are defined at the beginning of Section 2. In Section 2.1 we show that with high probability an LTF drawn from $\mathcal{D}_{no}$ is constant-far from monotone, and in Section 2.2 we show that unless $q = \Omega(n^{1/5}(\log n)^{-2/5})$, any deterministic q-query algorithm cannot distinguish between the two distributions with non-negligible success probability. The key technical ingredient in our proof of the latter is a lemma that adapts the Valiant–Valiant multidimensional CLT for Wasserstein distance to our context; we prove this lemma in Section 3. Finally in Section 4 we prove Theorem 2, showing that the same lower bound of $\tilde{\Omega}(n^{1/5})$ also applies to the query complexity of testers for monotonicity of functions $f : [m]^n \to \{0,1\}$ over general hypergrid domains; we do so via a reduction to the $m = 2$ case (Theorem 1).

Our algorithmic result is established in Section 5. In Section 5.1 we describe two useful distributions over comparable pairs $(\boldsymbol{x}, \boldsymbol{y})$ from the middle layers of $\{0,1\}^n$ and bound the probability

[3] Here the "middle layers" of $\boldsymbol{p}$ are the points on the path that have $n/2 \pm O_\varepsilon(\sqrt{n})$ many coordinates which are 1; intuitively, at most an ε-fraction of all points in $\{0,1\}^n$ lie outside these "middle layers" of the hypercube. We note that the above description is a slight simplification of the actual [CS13a] path tester, omitting some details which are not necessary at this stage of our description.

of having both points landing in a fixed set A of size $\sigma 2^n$. Then in Section 5.2 we define the *score* of a point x with respect to a set A of points, and use the result of Section 5.1 to lower bound the sum of score(x, A) over all points $x \in A$. We present our modified path tester as well as the analysis of its success probability in Section 5.3. Finally in Section 5.4 we combine this tester and the dichotomy theorem of [CS13a] to obtain our improved upper bound.

1.3 Preliminaries

All probabilities and expectations are with respect to the uniform distribution unless otherwise stated; we will use boldface letters (e.g. $\boldsymbol{x}$ and $\mathbf{X}$) to denote random variables. For a $q \times n$ matrix $Q \in \mathbb{R}^{q \times n}$, we write $Q_{i*} \in \mathbb{R}^n$ to denote its i-th row, $Q_{*j} \in \mathbb{R}^q$ its j-th column, and $Q_{i,j} \in \mathbb{R}$ its entry in the i-th column and j-th row. We write $\prec$ to denote the coordinate-wise partial order on $\{-1,1\}^n$, where $x \prec y$ iff $x_i \leq y_i$ for all $i \in [n]$ and $x \neq y$. We say that x and y are *comparable* if $x \prec y$, $y \prec x$, or $x = y$. Given two functions $f, g : \{-1,1\}^n \to \{-1,1\}$ we will use dist(f, g) to denote the (normalized Hamming) distance $\mathbf{Pr}_{\boldsymbol{x} \in \{-1,1\}^n}[f(\boldsymbol{x}) \neq g(\boldsymbol{x})]$ between f and g.

Recall that $f : \{-1,1\}^n \to \{-1,1\}$ is monotone if $f(x) \leq f(y)$ for all $x, y \in \{-1,1\}^n$ such that $x \prec y$. We say that f is *ε-close to monotone* if dist$(f, g) \leq \varepsilon$ for some monotone $g : \{-1,1\}^n \to \{-1,1\}$, and *$\varepsilon$-far from monotone* otherwise. A linear threshold function (LTF) over $\{-1,1\}^n$ is a function $f : \{-1,1\}^n \to \{-1,1\}$ that can be expressed as $f(x) = \text{sign}(w \cdot x - \theta)$ for some $w_1, \ldots, w_n$, $\theta \in \mathbb{R}$. Here sign $: \mathbb{R} \to \{-1,1\}$ is the sign function sign$(t) = 1$ if $t \geq 0$ and sign$(t) = -1$ if $t < 0$. For $f(x) = \text{sign}(w \cdot x - \theta)$, an LTF over $\{-1,1\}^n$, it is straightforward to verify that if $w_i \geq 0$ for all $i \in [n]$ then f is monotone.

We will need a few standard facts from probability theory:

Fact 1.1 (Gaussian anti-concentration)**.** *Let $\mathcal{G}$ be a Gaussian with variance σ^2. Then for all $\varepsilon > 0$ it holds that* $\sup_{\theta \in \mathbb{R}} \left\{ \mathbf{Pr}\left[|\mathcal{G} - \theta| \leq \varepsilon\sigma\right] \right\} \leq \varepsilon$.

Fact 1.2 (Gaussian concentration)**.** *Let $\mathcal{G}$ be a Gaussian with mean 0 and variance σ^2. Then for all $0 < a < 1$ it holds that* $\mathbf{Pr}\left[\mathcal{G} \in [0, a\sigma]\right] = \Omega(a)$.

Theorem 4 (Berry–Esséen)**.** *Let $\mathbf{S} = \mathbf{X}_1 + \cdots + \mathbf{X}_n$ where $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are independent real-valued random variables with $\mathbf{E}[\mathbf{X}_j] = \mu_j$ and $\mathbf{Var}[\mathbf{X}_j] = \sigma_j^2$, and suppose that $|\mathbf{X}_j - \mathbf{E}[\mathbf{X}_j]| \leq \tau$ with probability 1 for all $j \in [n]$. Let $\mathcal{G}$ be a Gaussian with mean $\sum_{j=1}^n \mu_j$ and variance $\sum_{j=1}^n \sigma_j^2$, matching those of $\mathbf{S}$. Then for all $\theta \in \mathbb{R}$, we have*

$$\left| \mathbf{Pr}[\mathbf{S} \leq \theta] - \mathbf{Pr}[\mathcal{G} \leq \theta] \right| \leq \frac{O(\tau)}{\left(\sum_{j=1}^n \sigma_j^2\right)^{1/2}}.$$

Fact 1.3. *For all $c > 0$ there exists an $\varepsilon = \varepsilon(c) \in (0, 1]$ such that the following holds. For all even (resp. odd) n,*

$$\mathop{\mathbf{Pr}}_{\boldsymbol{x} \in \{-1,1\}^n} \left[\sum_{i=1}^n \boldsymbol{x}_i = k \right] \geq \frac{\varepsilon}{\sqrt{n}} \quad \textit{for all even (resp. odd) integers } k \in [-c\sqrt{n}, c\sqrt{n}].$$

Finally we recall a few basic facts from the Fourier analysis over the hypercube which we require (for a comprehensive treatment of this topic see [O'D14]). Every function $f : \{-1,1\}^n \to \mathbb{R}$ can

be uniquely expressed as a multilinear polynomial

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i \quad \text{where } \widehat{f}(S) := \mathop{\mathbf{E}}_{\boldsymbol{x} \in \{-1,1\}^n} \left[f(\boldsymbol{x}) \prod_{i \in S} \boldsymbol{x}_i \right],$$

known as the *Fourier transform* of f. The numbers $\widehat{f}(S) \in \mathbb{R}$ are the *Fourier coefficients* of f; with a slight abuse of notation we will write $\widehat{f}(i)$ instead of $\widehat{f}(\{i\})$ for the degree-1 Fourier coefficients.

Fact 1.4 (Parseval's identity)**.** *Let $f : \{-1,1\}^n \to \mathbb{R}$. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{x} \in \{-1,1\}^n} [f(\boldsymbol{x})] = \sum_{S \subseteq [n]} \widehat{f}(S)^2.$$

For $i \in [n]$, the *influence of coordinate i on f*, denoted $\mathbf{Inf}_i[f]$, is the probability

$$\mathbf{Inf}_i[f] := \mathop{\mathbf{Pr}}_{\boldsymbol{x} \in \{-1,1\}^n} \left[f(\boldsymbol{x}) \neq f(\boldsymbol{x}^{\oplus i}) \right],$$

where $\boldsymbol{x}^{\oplus i}$ denotes the string $\boldsymbol{x}$ with its i-th coordinate flipped. The following fact relates the influences of an LTF to its degree-1 Fourier coefficients:

Fact 1.5. *Let $f(x) = \mathrm{sign}(w_1 x_1 + \cdots + w_n x_n - \theta)$ be an LTF over $\{-1,1\}^n$. Then for all $i \in [n]$, $\mathbf{Inf}_i[f] = \widehat{f}(i)$ if $w_i \geq 0$ and $\mathbf{Inf}_i[f] = -\widehat{f}(i)$ if $w_i < 0$.*

2 The lower bound: Proof of Theorem 2

Let $\mathcal{D}_{yes}$ be the following distribution over monotone LTFs on $\{-1,1\}^n$: a draw $\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}$ is $\boldsymbol{f}_{yes}(x) = \mathrm{sign}(\boldsymbol{\sigma}_1 x_1 + \cdots + \boldsymbol{\sigma}_n x_n)$, where each $\boldsymbol{\sigma}_i$ is independently and uniformly chosen from $\{1,3\}$. The distribution $\mathcal{D}_{no}$ is similarly a distribution over LTFs $\boldsymbol{f}_{no}(x) = \mathrm{sign}(\boldsymbol{\nu}_1 x_1 + \cdots + \boldsymbol{\nu}_n x_n)$, but each $\boldsymbol{\nu}_i$ is independently chosen to be -1 with probability $1/10$, and $7/3$ with probability $9/10$. The following two propositions along with a standard application of Yao's minimax principle [Yao77] yield Theorem 2:

Proposition 2.1. *There exists a universal positive constant $\varepsilon_0 > 0$ such that with probability $1 - o_n(1)$, a random LTF $\boldsymbol{f}_{no} \sim \mathcal{D}_{no}$ satisfies $\mathrm{dist}(\boldsymbol{f}_{no}, g) > \varepsilon_0$ for all monotone Boolean functions $g : \{-1,1\}^n \to \{-1,1\}$.*

Proposition 2.2. *Let $\mathcal{T}$ be any deterministic non-adaptive two-sided q-query algorithm for testing whether a black-box Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ is monotone. Then*

$$\left| \mathop{\mathbf{Pr}}_{\boldsymbol{f}_{yes} \sim \mathcal{D}_{yes}} \left[\mathcal{T} \text{ outputs } \mathsf{Yes} \text{ on } \boldsymbol{f}_{yes} \right] - \mathop{\mathbf{Pr}}_{\boldsymbol{f}_{no} \sim \mathcal{D}_{no}} \left[\mathcal{T} \text{ outputs } \mathsf{Yes} \text{ on } \boldsymbol{f}_{no} \right] \right| = O\left(\frac{q^{5/4} (\log n)^{1/2}}{n^{1/4}} \right). \quad (1)$$

We prove Proposition 2.1 in Section 2.1, followed by Proposition 2.2 in Section 2.2.

2.1 Proof of Proposition 2.1

By the Chernoff bound, with probability $1-o_n(1)$ a draw $\boldsymbol{f}_{no} = \mathrm{sign}(\boldsymbol{\nu}_1 x_1 + \cdots + \boldsymbol{\nu}_n x_n)$ from $\mathcal{D}_{no}$ satisfies

$$|\{i \in [n] \colon \boldsymbol{\nu}_i = -1\}| \in \left[0.1n - \sqrt{n \log n}, 0.1n + \sqrt{n \log n}\right].$$

We call any such LTF *nice*, and we will argue that all nice LTFs are constant-far from monotonicity. For the remainder of this proof let f be a nice LTF, which we may without loss of generality express as $f(x) = \mathrm{sign}(\ell(x))$ where

$$\ell(x) := -(x_1 + \cdots + x_m) + \tfrac{7}{3} \cdot (x_{m+1} + \cdots + x_n)$$

and $m \in \left[0.1n - \sqrt{n \log n}, 0.1n + \sqrt{n \log n}\right]$. We assume that m is odd, noting that the case when m is even follows via an identical argument. We first claim that $\mathbf{Inf}_i[f] = \Omega(1/\sqrt{n})$ for all $i \in [m]$; by symmetry it suffices to show this for $i = 1$. Define $\ell'(x) := -(x_2 + \ldots + x_m) + \frac{7}{3}(x_{m+1} + \cdots + x_n)$ and note that $f(x) \neq f(x^{\oplus 1})$ if and only if $\ell'(x) \in [-1, 1)$. Applying Fact 1.3 twice, we have

$$\mathop{\mathbf{Pr}}_{\boldsymbol{x} \in \{-1,1\}^n}\left[\tfrac{7}{3}(\boldsymbol{x}_{m+1} + \cdots + \boldsymbol{x}_n) \in [-\sqrt{n}, \sqrt{n}]\right] = \Omega(1)$$

and

$$\mathop{\mathbf{Pr}}_{\boldsymbol{x} \in \{-1,1\}^n}[\boldsymbol{x}_2 + \cdots + \boldsymbol{x}_m = k] = \Omega(1/\sqrt{n}) \text{ for all even integers } k \in [-\sqrt{n} - 1, \sqrt{n} + 1],$$

and therefore indeed,

$$\mathbf{Inf}_1[f] = \mathop{\mathbf{Pr}}_{\boldsymbol{x} \in \{-1,1\}^n}[\ell'(\boldsymbol{x}) \in [-1, 1)] = \Omega(1/\sqrt{n}).$$

Since $\mathbf{Inf}_i[f] = \Omega(1/\sqrt{n})$ for all $i \in [m]$, by Fact 1.5 we have that $\widehat{f}(i) = -\Omega(1/\sqrt{n})$ for all $i \in [m]$. Hence for all monotone Boolean functions g, we have

$$\begin{aligned} 4 \cdot \mathrm{dist}(f, g) \;=\; \mathop{\mathbf{E}}_{\boldsymbol{x} \in \{-1,1\}^n}\left[(f(\boldsymbol{x}) - g(\boldsymbol{x}))^2\right] \;&=\; \sum_{S \subseteq [n]} \left(\widehat{f}(S) - \widehat{g}(S)\right)^2 \\ &\geq\; \sum_{i=1}^{m} \left(\widehat{f}(i) - \widehat{g}(i)\right)^2 \;=\; m \cdot \Omega(1/n) \;=\; \Omega(1). \end{aligned}$$

Here the second equality is by Parvseval's identity; the penultimate equality uses the fact that $\widehat{g}(i) \geq 0$ for all $i \in [n]$, which in turn holds since g is a monotone Boolean function. This completes the proof of Proposition 2.1.

2.2 Proof of Proposition 2.2

Let $\mathcal{T}$ be any deterministic non-adaptive q-query tester, and view its q queries as a $q \times n$ matrix $Q \in \{-1, 1\}^{q \times n}$. Following the terminology of [BO10], we define a "Response Vector" random variable $\mathbf{R}_{yes} \in \{-1, 1\}^q$ which is obtained by drawing $\boldsymbol{f}_{yes} = \mathrm{sign}(\boldsymbol{\sigma}_1 x_1 + \cdots + \boldsymbol{\sigma}_n x_n)$ from $\mathcal{D}_{yes}$ and setting the i-th coordinate of $\mathbf{R}_{yes}$ to be

$$\boldsymbol{f}_{yes}(Q_{i*}) = \mathrm{sign}(\boldsymbol{\sigma}_1 Q_{i,1} + \cdots + \boldsymbol{\sigma}_n Q_{i,n}),$$

and similarly $\mathbf{R}_{no} \in \{-1,1\}^q$ which is obtained by drawing $\boldsymbol{f}_{no} \sim \mathcal{D}_{no}$ and setting the i-th coordinate of $\mathbf{R}_{no}$ to be $\boldsymbol{f}_{no}(Q_{i*})$. By the definition of total variation distance, the left-hand side of (1) is upper bounded by $d_{\mathrm{TV}}(\mathbf{R}_{yes}, \mathbf{R}_{no})$, and hence we can prove Proposition 2.2 by showing that $d_{\mathrm{TV}}(\mathbf{R}_{yes}, \mathbf{R}_{no}) = O(q^{5/4}(\log n)^{1/2}/n^{1/4})$.

Let $\mathbf{S} \in \mathbb{R}^q$ be the random column vector $Q\boldsymbol{\sigma}$ where $\boldsymbol{\sigma}$ is uniform over $\{1,3\}^n$, and $\mathbf{T} \in \mathbb{R}^q$ be the random column vector $Q\boldsymbol{\nu}$ where $\boldsymbol{\nu}$ is drawn from the product distribution over $\{-1, 7/3\}^n$ where $\mathbf{Pr}[\boldsymbol{\nu}_i = -1] = 1/10$ for all $i \in [n]$. The Response Vector $\mathbf{R}_{yes}$ is determined by the orthant of $\mathbb{R}^q$ in which $\mathbf{S}$ lies (as each coordinate of $\mathbf{R}_{yes}$ is simply the sign of the respective coordinate of $\mathbf{S}$), and likewise $\mathbf{R}_{no}$ by the orthant of $\mathbb{R}^q$ in which $\mathbf{T}$ lies. Therefore it suffices for us to prove the following lemma:

Lemma 2.3. *Let $\mathbf{S}, \mathbf{T} \in \mathbb{R}^q$ be defined as above. Then for any union $\mathcal{O}$ of orthants in $\mathbb{R}^q$,*

$$\left|\mathbf{Pr}[\mathbf{S} \in \mathcal{O}] - \mathbf{Pr}[\mathbf{T} \in \mathcal{O}]\right| = O\left(\frac{q^{5/4}(\log n)^{1/2}}{n^{1/4}}\right).$$

We will need the following multidimensional Berry–Esséen theorem, the proof of which we defer to Section 3.

Theorem 5. *Let $\mathbf{S} = \mathbf{X}^{(1)} + \cdots + \mathbf{X}^{(n)}$ where $\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(n)}$ are independent $\mathbb{R}^q$-valued random variables, and suppose that $\left|\mathbf{X}_i^{(j)} - \mathbf{E}[\mathbf{X}_i^{(j)}]\right| \le \tau$ with probability 1 for all $i \in [q]$ and $j \in [n]$. Let $\mathcal{G}$ be the q-dimensional Gaussian with the same mean and covariance matrix as $\mathbf{S}$. Let $\mathcal{O}$ be a union of orthants in $\mathbb{R}^q$. Then for all $r > 0$,*

$$\left|\mathbf{Pr}[\mathbf{S} \in \mathcal{O}] - \mathbf{Pr}[\mathcal{G} \in \mathcal{O}]\right| = O\left(\frac{\tau q^{3/2}\log n}{r} + \sum_{i=1}^{q} \frac{r+\tau}{\left(\sum_{j=1}^{n} \mathbf{Var}\left[\mathbf{X}_i^{(j)}\right]\right)^{1/2}}\right).$$

Proof of Lemma 2.3 assuming Theorem 5. We begin by writing $\mathbf{S} = \mathbf{X}^{(1)} + \cdots + \mathbf{X}^{(n)}$, where $\mathbf{X}^{(j)} = \boldsymbol{\sigma}_j \cdot Q_{*j}$ and $\boldsymbol{\sigma}_j$ is uniform over $\{1,3\}$; *i.e.* each $\mathbf{X}^{(j)}$ is independently Q_{*j} with probability $1/2$ and $3 \cdot Q_{*j}$ with probability $1/2$. Likewise we may express $\mathbf{T} = \mathbf{Y}^{(1)} + \cdots + \mathbf{Y}^{(n)}$, where $\mathbf{Y}^{(j)} = \boldsymbol{\nu}_j \cdot Q_{*j}$ and $\boldsymbol{\nu}_j$ is -1 with probability $1/10$ and $7/3$ with probability $9/10$. We claim that the $\mathbf{X}^{(j)}$'s and $\mathbf{Y}^{(j)}$'s have matching means and covariance matrices; it suffices to check this for $\mathbf{X}^{(1)}$ and $\mathbf{Y}^{(1)}$. For means, we see that indeed

$$\begin{aligned}
\mathbf{E}\left[\mathbf{X}^{(1)}\right] &= \mathbf{E}[\boldsymbol{\sigma}_1] \cdot Q_{*1} = (\tfrac{1}{2} + \tfrac{3}{2}) \cdot Q_{*1} = 2 \cdot Q_{*1} \\
\mathbf{E}\left[\mathbf{Y}^{(1)}\right] &= \mathbf{E}[\boldsymbol{\nu}_1] \cdot Q_{*1} = (-\tfrac{1}{10} + \tfrac{9}{10}\tfrac{7}{3}) \cdot Q_{*1} = 2 \cdot Q_{*1}.
\end{aligned}$$

As for the covariance matrices, we let $i_1, i_2 \in [q]$ and calculate

$$\begin{aligned}
\mathbf{Cov}[\mathbf{X}^{(1)}]_{i_1,i_2} &= \mathbf{E}\left[\left(\mathbf{X}_{i_1}^{(1)} - 2 \cdot Q_{i_1,1}\right)\left(\mathbf{X}_{i_2}^{(1)} - 2Q_{i_2,1}\right)\right] \\
&= \mathbf{E}\left[\mathbf{X}_{i_1}^{(1)} \cdot \mathbf{X}_{i_2}^{(1)}\right] - 2 \cdot Q_{i_2,1}\,\mathbf{E}\left[\mathbf{X}_{i_1}^{(1)}\right] - 2 \cdot Q_{i_1,1}\,\mathbf{E}\left[\mathbf{X}_{i_2}^{(1)}\right] + 4 \cdot Q_{i_1,1}Q_{i_2,1} \\
&= \mathbf{E}\left[\mathbf{X}_{i_1}^{(1)} \cdot \mathbf{X}_{i_2}^{(1)}\right] - 4 \cdot Q_{i_1,1}Q_{i_2,1} \\
&= \left(\mathbf{E}\left[\boldsymbol{\sigma}_1^2\right] - 4\right) \cdot Q_{i_1,1}Q_{i_2,1} \;=\; (\tfrac{1}{2} + \tfrac{9}{2} - 4) \cdot Q_{i_1,1}Q_{i_2,1} \;= Q_{i_1,1}Q_{i_2,1}.
\end{aligned}$$

Similarly, the corresponding entry of $\mathbf{Cov}[\mathbf{Y}^{(1)}]$ is:

$$\begin{aligned}
\mathbf{Cov}[\mathbf{Y}^{(1)}]_{i_1,i_2} &= \mathbf{E}\left[\left(\mathbf{Y}^{(1)}_{i_1} - 2\cdot Q_{i_1,1}\right)\left(\mathbf{Y}^{(1)}_{i_2} - 2Q_{i_2,1}\right)\right] \\
&= \mathbf{E}\left[\mathbf{Y}^{(1)}_{i_1}\cdot\mathbf{Y}^{(1)}_{i_2}\right] - 2\cdot Q_{i_2,1}\,\mathbf{E}\left[\mathbf{Y}^{(1)}_{i_1}\right] - 2\cdot Q_{i_1,1}\,\mathbf{E}\left[\mathbf{Y}^{(1)}_{i_2}\right] + 4\cdot Q_{i_1,1}Q_{i_2,1} \\
&= \mathbf{E}\left[\mathbf{Y}^{(1)}_{i_1}\cdot\mathbf{Y}^{(1)}_{i_2}\right] - 4\cdot Q_{i_1,1}Q_{i_2,1} \\
&= \left(\mathbf{E}\left[\boldsymbol{\nu}_1^2\right] - 4\right)\cdot Q_{i_1,1}Q_{i_2,1} \;=\; (\tfrac{1}{10} + \tfrac{9}{10}\tfrac{49}{9} - 4)\cdot Q_{i_1,1}Q_{i_2,1} \;= Q_{i_1,1}Q_{i_2,1}.
\end{aligned}$$

Since the $\mathbf{X}^{(j)}$'s and $\mathbf{Y}^{(j)}$'s have matching means and covariance matrices, so do their sums $\mathbf{S}$ and $\mathbf{T}$, and so Theorem 5 gives a bound on the differences $|\mathbf{Pr}[\mathbf{S}\in\mathcal{O}] - \mathbf{Pr}[\mathcal{G}\in\mathcal{O}]|$ and $|\mathbf{Pr}[\mathbf{T}\in\mathcal{O}] - \mathbf{Pr}[\mathcal{G}\in\mathcal{O}]|$ for the same q-dimensional Gaussian $\mathcal{G}$. Recalling that $\mathbf{X}^{(j)}_i = \boldsymbol{\sigma}_j\cdot Q_{i,j}$ where $Q_{i,j}\in\{-1,1\}^n$, we have that $\mathbf{Var}[\mathbf{X}^{(j)}_i] = 1$, and likewise $\mathbf{Var}[\mathbf{Y}^{(j)}_i] = 1$. Therefore, two applications of Theorem 5 with $\tau := O(1)$ along with the triangle inequality yields the bound

$$\left|\,\mathbf{Pr}[\mathcal{S}\in\mathcal{O}] - \mathbf{Pr}[\mathcal{G}\in\mathcal{O}]\right| = O\left(\frac{q^{3/2}\log n}{r} + \frac{q(r+\tau)}{\sqrt{n}}\right)$$

for all $r > 0$. Choosing $r := (qn)^{1/4}(\log n)^{1/2}$ completes the proof. □

3 Multidimensional Berry–Esséen via the Valiant–Valiant CLT

In this section we prove Theorem 5 by adapting a recent multidimensional CLT of Valiant and Valiant [VV11] which bounds the *Wasserstein distance* between a sum of independent vector-valued random variables and a multidimensional Gaussian.

Definition 6 (Wasserstein distance)**.** The Wasserstein distance between two $\mathbb{R}^q$-valued random variables $\mathbf{S}$ and $\mathbf{T}$, denoted $d_W(\mathbf{S},\mathbf{T})$, is defined to be:

$$d_W(\mathbf{S},\mathbf{T}) = \inf_{\mathcal{D}}\left\{\mathop{\mathbf{E}}_{\mathcal{D}}\left[\|\mathbf{U}-\mathbf{V}\|_2\right]\right\},$$

where the infimum is taken over all couplings $\mathcal{D}$ of $\mathbf{S}$ and $\mathbf{T}$, i.e., all joint distributions $\mathcal{D}$ of pairs of $\mathbb{R}^q$-valued random variables $(\mathbf{U},\mathbf{V})$ with marginals distributed according to $\mathbf{S}$ and $\mathbf{T}$ respectively.

Valiant and Valiant [VV11] recently used Stein's method to prove the following central limit theorem for Wasserstein distance:

Theorem 7 (Valiant–Valiant CLT)**.** *Let $\mathbf{S} = \mathbf{X}^{(1)}+\cdots+\mathbf{X}^{(n)}$ where $\mathbf{X}^{(1)},\ldots,\mathbf{X}^{(n)}$ are independent $\mathbb{R}^q$-valued random variables, and suppose $\left\|\mathbf{X}^{(j)} - \mathbf{E}\left[\mathbf{X}^{(j)}\right]\right\|_2 \le \beta$ with probability 1 for any $j\in[n]$. Then*

$$d_W(\mathbf{S},\mathcal{G}) \le O(\beta q\log n),$$

where $\mathcal{G}$ is the q-dimensional Gaussian with the same mean and covariance matrix as $\mathbf{S}$.

We recall Theorem 5:

Theorem 5. *Let $\mathbf{S} = \mathbf{X}^{(1)} + \cdots + \mathbf{X}^{(n)}$ where $\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(n)}$ are independent $\mathbb{R}^q$-valued random variables, and suppose that $\left|\mathbf{X}_i^{(j)} - \mathbf{E}\left[\mathbf{X}_i^{(j)}\right]\right| \le \tau$ with probability 1 for all $i \in [q]$ and $j \in [n]$. Let $\mathcal{G}$ be the q-dimensional Gaussian with the same mean and covariance matrix as $\mathbf{S}$. Let $\mathcal{O}$ be a union of orthants in $\mathbb{R}^q$. Then for all $r > 0$,*

$$\left|\mathbf{Pr}[\mathbf{S} \in \mathcal{O}] - \mathbf{Pr}[\mathcal{G} \in \mathcal{O}]\right| = O\left(\frac{\tau q^{3/2} \log n}{r} + \sum_{i=1}^{q} \frac{r + \tau}{\left(\sum_{j=1}^{n} \mathbf{Var}\left[\mathbf{X}_i^{(j)}\right]\right)^{1/2}}\right).$$

Proof. We define

$$W_r := \left\{x \in \mathbb{R}^q \colon |x_i| \le r \text{ for some } i \in [q]\right\}$$

to be the radius-r region around the orthant boundaries, and partition $\mathcal{O}$ into $\mathcal{O}_{bd} := \mathcal{O} \cap W_r$ (the points in $\mathcal{O}$ that lie close to the orthant boundaries) and $\mathcal{O}_{in} := \mathcal{O} \setminus W_r$ (the points that lie far away from the orthant boundaries). We have

$$\begin{aligned}\left|\mathbf{Pr}[\mathbf{S} \in \mathcal{O}] - \mathbf{Pr}[\mathcal{G} \in \mathcal{O}]\right| &= \left|(\mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{in}] + \mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{bd}]) - (\mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{in}] + \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{bd}])\right| \\ &\le \underbrace{\left|\mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{in}] - \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{in}]\right|}_{\Delta} + \underbrace{\mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{bd}] + \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{bd}]}_{\Gamma}. \quad (2)\end{aligned}$$

We bound the quantities Δ and Γ separately. For Γ, we have that

$$\begin{aligned}\Gamma &\le \sum_{i=1}^{q} \mathbf{Pr}\left[\mathbf{S}_i \in [-r, r]\right] + \mathbf{Pr}\left[\mathcal{G}_i \in [-r, r]\right] \quad (3) \\ &\le \sum_{i=1}^{q} 2\,\mathbf{Pr}\left[\mathcal{G}_i \in [-r, r]\right] + \left|\mathbf{Pr}\left[\mathbf{S}_i \in [-r, r]\right] - \mathbf{Pr}\left[\mathcal{G}_i \in [-r, r]\right]\right| \\ &\le \sum_{i=1}^{q} \frac{O(r)}{\left(\sum_{j=1}^{n} \mathbf{Var}\left[\mathbf{X}_i^{(j)}\right]\right)^{1/2}} + \frac{O(\tau)}{\left(\sum_{j=1}^{n} \mathbf{Var}\left[\mathbf{X}_i^{(j)}\right]\right)^{1/2}} = \sum_{i=1}^{q} \frac{O(r + \tau)}{\left(\sum_{j=1}^{n} \mathbf{Var}\left[\mathbf{X}_i^{(j)}\right]\right)^{1/2}} \quad (4)\end{aligned}$$

where (3) is a union bound over all q dimensions, and (4) uses Fact 1.1 (Gaussian anti-concentration), the fact that $\mathcal{G}_i$ is a Gaussian with variance $\sum_{j=1}^{n} \mathbf{Var}\left[\mathbf{X}_i^{(j)}\right]$, and Theorem 4 (Berry–Esséen).

For Δ, let us assume without loss of generality (a symmetrical argument works in the other case) that $\mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{in}] \ge \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{in}]$, so $\Delta = \mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{in}] - \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{in}]$. Let $\mathcal{D}$ be any coupling of $\mathbf{S}$ and $\mathcal{G}$, so $\mathcal{D}$ is the joint distribution of a pair $(\mathbf{U}, \mathbf{V})$ of $\mathbb{R}^q$-valued random variables with marginals distributed according to $\mathbf{S}$ and $\mathcal{G}$ respectively. Since

$$\int_{\mathcal{O}_{in}} \int_{\mathbb{R}^q} \mathcal{D}(u, v)\, dv\, du = \mathbf{Pr}[\mathbf{S} \in \mathcal{O}_{in}]$$

and

$$\int_{\mathcal{O}_{in}} \int_{\mathcal{O}_{in}} \mathcal{D}(u, v)\, dv\, du \le \int_{\mathbb{R}^q} \int_{\mathcal{O}_{in}} \mathcal{D}(u, v)\, dv\, du = \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{in}],$$

it follows that

$$\int_{\mathcal{O}_{in}} \int_{\mathbb{R}^q \setminus \mathcal{O}_{in}} \mathcal{D}(u, v)\, dv\, du = \int_{\mathcal{O}_{in}} \int_{\mathbb{R}^q} \mathcal{D}(u, v)\, dv\, du - \int_{\mathcal{O}_{in}} \int_{\mathcal{O}_{in}} \mathcal{D}(u, v)\, dv\, du \ge \Delta. \quad (5)$$

Next we define the quantities

$$\begin{aligned}\Delta_{near}(\mathcal{D}) &:= \int_{\mathcal{O}_{in}} \int_{\mathcal{O}_{bd}} \mathcal{D}(u,v)\,dv\,du \\ \Delta_{far}(\mathcal{D}) &:= \int_{\mathcal{O}_{in}} \int_{\mathbb{R}^q \setminus \mathcal{O}} \mathcal{D}(u,v)\,dv\,du.\end{aligned}$$

Note that $\Delta_{near}(\mathcal{D})$ and $\Delta_{far}(\mathcal{D})$ sum to the quantity on the left-hand side of (5), and so $\Delta_{near}(\mathcal{D}) + \Delta_{far}(\mathcal{D}) \geq \Delta$. (In words, since $\mathbf{S}$ places Δ more mass on $\mathcal{O}_{in}$ than $\mathcal{G}$ does, any scheme $\mathcal{D}$ of moving the mass of $\mathbf{S}$ to obtain $\mathcal{G}$ must move at least Δ amount from within $\mathcal{O}_{in}$ to outside it. $\Delta_{near}(\mathcal{D})$ is the amount moved from within $\mathcal{O}_{in}$ to $\mathcal{O}$'s boundary $\mathcal{O}_{bd}$, and $\Delta_{far}(\mathcal{D})$ is the rest, moved from within $\mathcal{O}_{in}$ to locations entirely out of $\mathcal{O}$.) Since $\|u - v\|_2 \geq r$ for any pair of points $u \in \mathcal{O}_{in}$ and $y \notin \mathcal{O}$, it follows that

$$d_W(\mathbf{S}, \mathcal{G}) \geq r \cdot \Delta_{far}(\mathcal{D}).$$

We consider two cases, depending on the relative magnitudes of $\Delta_{near}(\mathcal{D})$ and $\Delta_{far}(\mathcal{D})$. If $\Delta_{far}(\mathcal{D}) \geq \Delta_{near}(\mathcal{D})$, we first observe that for all $j \in [n]$ we have $\left\|\mathbf{X}^{(j)} - \mathbf{E}\left[\mathbf{X}^{(j)}\right]\right\|_2 \leq \tau\sqrt{q}$ with probability 1, since each of its q coordinates i satisfies $\left|\mathbf{X}_i^{(j)} - \mathbf{E}\left[\mathbf{X}_i^{(j)}\right]\right| \leq \tau$ with probability 1 by the assumption of the theorem. Therefore we may apply Theorem 7 (Valiant–Valiant CLT) with $\beta := \tau\sqrt{q}$ to get

$$r \cdot \frac{\Delta}{2} \leq r \cdot \Delta_{far}(\mathcal{D}) \leq d_W(\mathbf{S}, \mathcal{G}) = O(\tau q^{3/2} \log n)$$

and hence $\Delta = O((\tau q^{3/2} \log n)/r)$, which along with our upper bound on Γ completes the proof. If on the other hand $\Delta_{near}(\mathcal{D}) > \Delta_{far}(\mathcal{D})$, then

$$\frac{\Delta}{2} \leq \Delta_{near}(\mathcal{D}) \leq \int_{\mathbb{R}^q} \int_{\mathcal{O}_{bd}} \mathcal{D}(u,v)\,dv\,du = \mathbf{Pr}[\mathcal{G} \in \mathcal{O}_{bd}] \leq \Gamma,$$

and again our bound on Γ completes the proof. □

4 A lower bound for general hypergrid domains

In this section we prove Theorem 2, showing that for all $m \in \mathbb{N}$ essentially the same lower bound of $\tilde{\Omega}(n^{1/5})$ also applies to the query complexity of testers for monotonicity of functions $f : [m]^n \to \{-1, 1\}$, Boolean-valued functions over general hypergrid domains. The notions of monotonicity and distance to monotonicity of functions generalize to functions $f : [m]^n \to \{-1, 1\}$ the natural way: f is monotone if $f(x) \leq f(y)$ for all $x \prec y$, where $x \prec y$ iff $x_i \leq y_i$ for all $i \in [n]$ and $x \neq y$. We say that f is ε-close to monotone if $\mathbf{Pr}_{x \in [m]^n}[f(x) \neq g(x)] \leq \varepsilon$ for some monotone $g : [m]^n \to \{-1, 1\}$, and ε-far from monotone otherwise.

We prove Theorem 2 via a reduction to the $m = 2$ case (i.e. Theorem 1). The reduction is simpler for even m so for ease of exposition we assume below that m is even. In this case Theorem 2 is a direct consequence of Theorem 1 and the following proposition:

Proposition 4.1. *For all even $m \in \mathbb{N}$ the mapping*

$$\Phi : \left\{\textit{all functions } f : \{-1,1\}^n \to \{-1,1\}\right\} \to \left\{\textit{all functions } f : [m]^n \to \{-1,1\}\right\}$$

defined by (6) below satisfies the following two properties:

1. *If $f : \{-1,1\}^n \to \{-1,1\}$ is monotone then $\Phi[f]$ is monotone as well.*

2. *If $f : \{-1,1\}^n \to \{-1,1\}$ is ε-far from monotone then $\Phi[f]$ is ε-far from monotone as well.*

We will need the following characterization of distance to monotonicity.

Theorem 8 ([FLN+02] Lemma 4)**.** *For all $f : [m]^n \to \{-1,1\}$ and $\varepsilon > 0$, we have that f is ε-far from monotone if and only if there exists εm^n many pairwise disjoint ordered pairs of vertices $(x^i, y^i) \in [m]^n \times [m]^n$ such that $x^i \prec y^i$ and $f(x^i) > f(y^i)$. We will call each such pair a* violation with respect to f.

Proof of Proposition 4.1. For every $f : \{-1,1\}^n \to \{-1,1\}$, we define $\Phi[f] : [m]^n \to \{-1,1\}$ to be the function

$$\Phi[f](x_1, \ldots, x_n) := f\big(\mathbf{1}[x_1 > m/2], \ldots, \mathbf{1}[x_n > m/2]\big), \tag{6}$$

where we use $\mathbf{1}[\cdot]$ to denote the $\{\pm 1\}$-valued indicator where $\mathbf{1}[P] = 1$ if P is true, and -1 otherwise. (Note that $m/2$ is an integer by our assumption that m is even.)

It is straightforward to verify that $\Phi[f]$ is monotone if f is monotone, and so it remains to show that $\Phi[f]$ is ε-far from monotone if f is ε-far from monotone. Since f is ε-far from monotone, we have by Theorem 8 that there exist $\varepsilon 2^n$ many pairwise disjoint pairs $(x^i, y^i) \in \{-1,1\}^n \times \{-1,1\}^n$ that are violations with respect to f; we will exhibit εm^n many pairwise disjoint pairs in $[m]^n$ that are violations with respect to $\Phi[f]$, which along with another application of Theorem 8 completes the proof. Let $S : \{-1,1\} \to \big\{[m/2], \{(m/2)+1, \ldots, m\}\big\}$ be the set-valued function

$$S(b) = \begin{cases} [m/2] & \text{if } b = -1 \\ \{(m/2)+1, \ldots, m\} & \text{if } b = 1, \end{cases}$$

and by a slight abuse of notation, we also define

$$S(x) = S(x_1) \times \cdots \times S(x_n) \subseteq [m]^n$$

to be a function that maps points $x \in \{-1,1\}^n$ to subsets of $[m]^n$. Note that $|S(x)| = (m/2)^n$ for all $x \in \{-1,1\}^n$, and $S(x) \cap S(y) = \emptyset$ if $x \neq y$. Furthermore, $\Phi[f](x') = f(x)$ for all $x \in \{-1,1\}^n$ and $x' \in S(x)$. In words, S maps each 1-input of f to a set of $(m/2)^n$ many 1-inputs of $\Phi[f]$, and likewise each 0-input of f to a set of $(m/2)^n$ many 0-inputs of $\Phi[f]$.

For any pair $(x, y) \in \{-1,1\}^n \times \{-1,1\}^n$ that is a violation with respect to f, consider pairing the $(m/2)^n$ elements of $S(x)$ with the $(m/2)^n$ elements of $S(y)$ in the obvious way (i.e. each $a = (a_1, \ldots, a_n) \in S(x)$ is is paired with the unique element $b = (b_1, \ldots, b_n) \in S(y)$ that has $(a_i \mod m/2) = (b_i \mod m/2)$ for all i). Since $x \prec y$, it follows from the definition of S that every $x' \in S(x)$ is paired with $y' \in S(y)$ where $x' \prec y'$. Furthermore, as noted above $\Phi[f](x') = f(x) = 1$ whereas $\Phi[f](y') = f(y) = 0$, and so every pair $(x', y') \in S(x) \times S(y)$ is a violation with respect to $\Phi[f]$. Therefore each of the $\varepsilon 2^n$ many pairs $(x, y) \in \{-1,1\}^n \times \{-1,1\}^n$ that are violations with respect to f gives rise to $(m/2)^n$ many pairwise disjoint pairs $(x', y') \in S(x) \times S(y)$ that are violations with respect to $\Phi[f]$. Finally recalling that $S(x) \cap S(y) = \emptyset$ if $x \neq y$, we conclude that there are indeed $\varepsilon 2^n \cdot (m/2)^n = \varepsilon m^n$ many pairwise disjoint pairs that are violations with respect to $\Phi[f]$. This finishes the proof. ∎

Remark 9. The proof for odd m, deferred to Appendix A, is via a similar but more involved version of Proposition 4.1. In place of the simple indicator function $\mathbf{1}[x_i > m/2]$ (whose domain

is simply $[m]$), we now use an "almost-balanced" monotone function $h : [m]^k \to \{-1, 1\}$ where $k = \Theta(\log n)$ and h has some additional properties. The fact that $k = \Theta(\log n)$ incurs an additional logarithmic loss in the parameters but still results in a $\tilde{\Omega}(n^{1/5})$ lower bound.

5 The algorithm

Throughout the proof of our upper bound we will assume that $1/n \le \varepsilon \le 1/2$. Note that this is without loss of generality, since if $\varepsilon < 1/n$ then the edge tester alone succeeds with probability $\Omega(\varepsilon/n) = \Omega(\varepsilon^2)$, and if $\varepsilon > 1/2$ then every f is ε-close to one of the two constant functions, both of which are monotone.

For our upper bound it will be more convenient to view Boolean functions as mapping $\{0,1\}^n$ to $\{0,1\}$. For $x, y \in \{0,1\}^n$ we write $\|x\|_1$ to denote $\sum_{i=1}^n x_i$, the number of 1s in x, and $\|x - y\|_1$ to denote $|\{i \in [n] \colon x_i \neq y_i\}|$, the ℓ_1-distance between x and y. Given $1/n \le \varepsilon \le 1/2$, we fix

$$d(n, \varepsilon) := 2\left\lceil \sqrt{2n \ln(100/\varepsilon)} \right\rceil = O\big(\sqrt{n \ln(1/\varepsilon)}\big),$$

and will denote $d(n, \varepsilon)$ simply by d when the distance parameter ε is clear from the context. For each $i \in \{0, 1, \dots, n\}$ we let $L_i := \{x \in \{0,1\}^n \colon \|x\|_1 = i\}$ denote the i-th layer, and refer to

$$L_{\text{mid}} := \big\{x \in L_i \colon i \in [(n-d)/2, (n+d)/2]\big\}$$

as the middle layers of the hypercube. A standard Chernoff bound gives $|\{0,1\}^n \setminus L_{\text{mid}}| \le (\varepsilon/50)2^n$. Finally, by a "path" we always mean a directed path of $n+1$ adjacent vertices from 0^n up to 1^n.

5.1 Two useful distributions over comparable pairs

Let $\mathcal{D} = \mathcal{D}_{n,\varepsilon}$ denote the following distribution over comparable pairs $(\boldsymbol{x}, \boldsymbol{y}) \in L_{\text{mid}} \times L_{\text{mid}}$:

1. First pick a path $\boldsymbol{p}$ uniformly from the collection of all paths going from 0^n to 1^n.
2. Pick $\boldsymbol{x}$ and $\boldsymbol{y}$ independently and uniformly from $\boldsymbol{p}_{\text{mid}} := \{z \in \boldsymbol{p} \colon z \in L_{\text{mid}}\}$.

This distribution is a slight variant of the one induced by the [CS13a] path tester, which takes a parameter σ as input and disallows pairs (x, y) for which $\|x - y\|_1$ is too small relative to σ. Our new tester will *not* sample from $\mathcal{D}$ (see Section 5.3), but we will use $\mathcal{D}$ in our analysis. We remark here that $\boldsymbol{x} = \boldsymbol{y}$ with positive probability under $\mathcal{D}$.

If $\boldsymbol{x}, \boldsymbol{y}$ were chosen independently and uniformly from $\{0,1\}^n$, then the probability that they both land in a fixed set A of $\sigma 2^n$ points, for some $\sigma \in (0,1)$, would be σ^2. The following lemma states that the probability is not much lower for a pair drawn from $\mathcal{D}$:

Lemma 5.1. *Let $A \subseteq L_{\text{mid}}$ be a set of $\sigma 2^n$ points. Then $\mathbf{Pr}_{(\boldsymbol{x},\boldsymbol{y}) \leftarrow \mathcal{D}}[\boldsymbol{x}, \boldsymbol{y} \in A] = \Omega\big(\sigma^2 \ln^{-1}(1/\varepsilon)\big)$.*

Proof. Applying Jensen's inequality, we have

$$\begin{aligned}\Pr_{(\boldsymbol{x},\boldsymbol{y}) \leftarrow \mathcal{D}}[\boldsymbol{x}, \boldsymbol{y} \in A] &= \mathop{\mathbf{E}}_{\boldsymbol{p}}\Big[\Pr_{\boldsymbol{x},\boldsymbol{y} \in \boldsymbol{p}}[\boldsymbol{x}, \boldsymbol{y} \in A]\Big] \\ &= \mathop{\mathbf{E}}_{\boldsymbol{p}}\left[\left(\frac{|\boldsymbol{p}_{\text{mid}} \cap A|}{|\boldsymbol{p}_{\text{mid}}|}\right)^2\right] = \Omega\left(\frac{1}{n \ln(1/\varepsilon)}\right) \cdot \mathop{\mathbf{E}}_{\boldsymbol{p}}\big[|\boldsymbol{p}_{\text{mid}} \cap A|\big]^2,\end{aligned}$$

and so it suffices to lower bound $\mathbf{E}_{\boldsymbol{p}}[|\boldsymbol{p}_{\text{mid}} \cap A|]$ by $\Omega(\sigma\sqrt{n})$. This is exactly Claim 2.2.1 of [CS13a]; we repeat the calculation here for the sake of completeness:

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{p}}\left[|\boldsymbol{p}_{\text{mid}} \cap A|\right] &= \mathop{\mathbf{E}}_{\boldsymbol{p}}\left[\sum_{i=\frac{1}{2}(n-d)}^{\frac{1}{2}(n+d)} \mathbf{1}\left[(\boldsymbol{p}_{\text{mid}} \cap L_i) \subseteq A\right]\right] \\
&= \sum_{i=\frac{1}{2}(n-d)}^{\frac{1}{2}(n+d)} \mathop{\mathbf{E}}_{\boldsymbol{p}}\left[\mathbf{1}\left[(\boldsymbol{p}_{\text{mid}} \cap L_i) \subseteq A\right]\right] \\
&= \sum_{i=\frac{1}{2}(n-d)}^{\frac{1}{2}(n+d)} \frac{|A \cap L_i|}{|L_i|} \qquad (7) \\
&\geq \frac{\sqrt{n}}{2^n} \sum_{i=\frac{1}{2}(n-d)}^{\frac{1}{2}(n+d)} |A \cap L_i| = \frac{|A|\sqrt{n}}{2^n} = \sigma\sqrt{n}, \qquad (8)
\end{aligned}$$

where we use $\mathbf{1}[\cdot]$ to denote the $\{0,1\}$-valued indicator where $\mathbf{1}[P] = 1$ if P is true, and 0 otherwise. Here (7) uses the fact that a uniformly random path $\boldsymbol{p}$ from 0^n to 1^n contains a uniformly random point in layer L_i, and (8) holds since $|L_i| \leq 2^n/\sqrt{n}$ for all i. □

We will need a numerical lemma concerning the ratio of binomial coefficients.

Lemma 5.2. *Let $\varepsilon \geq 1/n$, and $a, b \in [(n-d)/2, (n+d)/2]$ be integers where $a > b$. Then*

$$\binom{a}{a-b} \Big/ \binom{n-b}{a-b} = O(1/\varepsilon^4) \quad \text{and} \quad \binom{n}{n/2} \Big/ \binom{n}{a} = O(1/\varepsilon^4).$$

Proof. We prove the first equation and the second equation is similar. By a routine calculation we verify that the first ratio is maximized when $a = (n+d)/2$ and $b = n/2$, and so

$$\frac{\binom{a}{a-b}}{\binom{n-b}{a-b}} \leq \frac{\frac{1}{2}(n+d)}{\frac{n}{2}} \cdot \frac{\frac{1}{2}(n+d)-1}{\frac{n}{2}-1} \cdots \frac{\frac{n}{2}+1}{\frac{1}{2}(n-d)+1} \leq \exp\left(\sum_{i=0}^{(d/2)-1} \frac{d/2}{(n/2)-i}\right),$$

where we used $(1+t) \leq e^t$ for $t \in \mathbb{R}$. The lemma follows from the definition of d and $\varepsilon \geq 1/n$. □

For our analysis, the following distribution $\mathcal{D}' = \mathcal{D}'_{n,\varepsilon}$ over comparable pairs $(\boldsymbol{x}, \boldsymbol{y}) \in L_{\text{mid}} \times L_{\text{mid}}$ in the middle layers comes in handy:

1. First pick a point $\boldsymbol{x}$ uniformly at random from L_{mid}.
2. Then pick a path $\boldsymbol{p}$ uniformly from the collection of all paths going through 0^n, $\boldsymbol{x}$, and 1^n.
3. Pick $\boldsymbol{y}$ uniformly from $\boldsymbol{p}_{\text{mid}} := \{z \in \boldsymbol{p} \colon z \in L_{\text{mid}}\}$.

We note that $\mathcal{D}'$ is not the same as $\mathcal{D}$, since picking a uniformly random $\boldsymbol{x}$ from the middle layers of a uniformly random path $\boldsymbol{p}$ does not induce a uniform distribution over L_{mid}; however, Lemma 5.2 allows us to switch between these essentially-equivalent distributions at the cost of a $O(1/\varepsilon^4)$ factor. (On the other hand the conditional distributions $\mathcal{D}_{\boldsymbol{x}=x}$ and $\mathcal{D}'_{\boldsymbol{x}=x}$ on $\boldsymbol{y}$ are the same for all possible outcomes $x \in L_{\text{mid}}$ of $\boldsymbol{x}$.)

We get the following corollary from Lemmas 5.1 and 5.2:

Corollary 5.3. *Let $A \subseteq L_{\mathrm{mid}}$ be a set of $\sigma 2^n$ points. Then*

$$\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}[\boldsymbol{x},\boldsymbol{y} \in A] = \Omega\left(\sigma^2\varepsilon^4 \ln^{-1}(1/\varepsilon)\right).$$

Proof. It is clear from the definition of $\mathcal{D}, \mathcal{D}'$ that the conditional distribution of $\boldsymbol{y}$ induced from $\mathcal{D}$ by conditioning on a particular outcome of $\boldsymbol{x}$ is the same as that induced from $\mathcal{D}'$ under the same conditioning. It follows from the second part of Lemma 5.2 that for any $x \in L_{\mathrm{mid}}$ we have

$$\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}[\boldsymbol{x} = x] = \Omega(\varepsilon^4) \cdot \mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}}[\boldsymbol{x} = x].$$

As a result, we have for every comparable pair (x, y) in the middle layers

$$\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}[(\boldsymbol{x},\boldsymbol{y}) = (x, y)] = \Omega(\varepsilon^4) \cdot \mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}}[(\boldsymbol{x},\boldsymbol{y}) = (x, y)].$$

The claim then follows from Lemma 5.1. ∎

5.2 Density and score

We need the following definition to give a more detailed analysis on the consequence of Corollary 5.3, which is key to the analysis of our monotonicity tester described in Section 5.3.

Definition 10 (density and score)**.** Let $A \subseteq \{0,1\}^n$. For all $x \in \{0,1\}^n$ and $k \in \{0, 1, \ldots, n\}$, we define the following quantities:

$$\mathrm{dens}^{\downarrow}_k(x, A) := \mathop{\mathbf{Pr}}_{\substack{\boldsymbol{y}\preceq x \\ \|\boldsymbol{y}-x\|_1=k}}[\boldsymbol{y} \in A] \text{ if } k \le \|x\|_1, \text{ and } \mathrm{dens}^{\downarrow}_k(x, A) := 0 \text{ otherwise,}$$

and similarly

$$\mathrm{dens}^{\uparrow}_k(x, A) := \mathop{\mathbf{Pr}}_{\substack{\boldsymbol{y}\succeq x \\ \|\boldsymbol{y}-x\|_1=k}}[\boldsymbol{y} \in A] \text{ if } k \le n - \|x\|_1, \text{ and } \mathrm{dens}^{\uparrow}_k(x, A) := 0 \text{ otherwise.}$$

We also define

$$\mathrm{score}^{\downarrow}(x, A) := \sum_{k=0}^{n} \mathrm{dens}^{\downarrow}_k(x, A) \quad \text{and} \quad \mathrm{score}^{\uparrow}(x, A) := \sum_{k=1}^{n} \mathrm{dens}^{\uparrow}_k(x, A),$$

and refer to $\mathrm{score}^{\downarrow}(x, A)$ as the *downward A-score of x* and $\mathrm{score}^{\uparrow}(x, A)$ as its *upward A-score*.

We point out the asymmetry between the definitions of $\mathrm{score}^{\downarrow}(x, A)$ and $\mathrm{score}^{\uparrow}(x, A)$: the first is summed over k starting at 0, whereas the second is summed over k starting at 1. (Note that $\mathrm{dens}^{\downarrow}_0(x, A) = \mathrm{dens}^{\uparrow}_0(x, A) = \mathbf{1}[x \in A]$.) We will need the fact that both the upward and downward A-scores of any $x \in \{0,1\}^n$ are at most $d = d(n, \varepsilon)$ when $A \subseteq L_{\mathrm{mid}}$.

The following lemma relates the distribution $\mathcal{D}'$ (more precisely, the distribution over $\boldsymbol{y}$ that is induced by conditioning on a particular outcome of $\boldsymbol{x}$) to the notion of score:

Lemma 5.4. *Let $A \subseteq L_{\mathrm{mid}}$ be a set of $\sigma 2^n$ points and fix $x^* \in L_{\mathrm{mid}}$. Then*

$$\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}\left[\boldsymbol{y} \in A \mid \boldsymbol{x} = x^*\right] = \frac{1}{\Theta(\sqrt{n \ln(1/\varepsilon)})}\left(\mathrm{score}^{\downarrow}(x^*, A) + \mathrm{score}^{\uparrow}(x^*, A)\right).$$

Proof. This holds since

$$\begin{aligned}
\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}\left[\boldsymbol{y}\in A \mid \boldsymbol{x}=x^*\right] &= \mathop{\mathbf{E}}_{\boldsymbol{p}\ni x^*}\left[\frac{|\boldsymbol{p}_{\text{mid}}\cap A|}{|\boldsymbol{p}_{\text{mid}}|}\right] \\
&= \frac{1}{\Theta(d)}\cdot\mathop{\mathbf{E}}_{\boldsymbol{p}\ni x^*}\left[|\boldsymbol{p}_{\text{mid}}\cap A|\right] \\
&= \frac{1}{\Theta(d)}\left(\sum_{k\geq 0}\left(\mathop{\mathbf{E}}_{\substack{\boldsymbol{y}\preceq x^*\\ \|\boldsymbol{y}-x^*\|_1=k}}\left[\mathbf{1}[\boldsymbol{y}\in A]\right]\right)+\sum_{k\geq 1}\left(\mathop{\mathbf{E}}_{\substack{\boldsymbol{y}\succ x^*\\ \|\boldsymbol{y}-x^*\|_1=k}}\left[\mathbf{1}[\boldsymbol{y}\in A]\right]\right)\right) \\
&= \frac{1}{\Theta(d)}\left(\sum_{k\geq 0}\left(\mathop{\mathbf{Pr}}_{\substack{\boldsymbol{y}\preceq x^*\\ \|\boldsymbol{y}-x^*\|_1=k}}\left[\boldsymbol{y}\in A\right]\right)+\sum_{k\geq 1}\left(\mathop{\mathbf{Pr}}_{\substack{\boldsymbol{y}\succ x^*\\ \|\boldsymbol{y}-x^*\|_1=k}}\left[\boldsymbol{y}\in A\right]\right)\right) \\
&= \frac{1}{\Theta(d)}\left(\text{score}^{\downarrow}(x^*,A)+\text{score}^{\uparrow}(x^*,A)\right).
\end{aligned}$$

□

We use the previous two lemmas to lower bound the expected downward A-score of an $\boldsymbol{x}$ drawn uniformly at random from A:

Lemma 5.5. *Let $\varepsilon\geq 1/n$ and $A\subseteq L_{\text{mid}}$ be a set of $\sigma 2^n$ points. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{x}\in A}\left[\text{score}^{\downarrow}(\boldsymbol{x},A)\right]=\Omega\left(\frac{\varepsilon^8\sigma\sqrt{n}}{\sqrt{\ln(1/\varepsilon)}}\right).$$

Proof. We begin with the claim that

$$\mathop{\mathbf{E}}_{\boldsymbol{x}\in A}\left[\text{score}^{\downarrow}(\boldsymbol{x},A)\right]\geq\Omega(\varepsilon^4)\mathop{\mathbf{E}}_{\boldsymbol{x}\in A}\left[\text{score}^{\uparrow}(\boldsymbol{x},A)\right]+1, \tag{9}$$

where the $+1$ is due to $\text{dens}_0^{\downarrow}(\boldsymbol{x},A)=1$. To see (9), we rewrite the LHS of the inequality as follows:

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{x}\in A}\left[\text{score}^{\downarrow}(\boldsymbol{x},A)\right]-1 &= \frac{1}{\sigma 2^n}\sum_{x\in A}\sum_{k\geq 1}\sum_{\substack{y\prec x\\ \|y-x\|_1=k}}\frac{\mathbf{1}[y\in A]}{\binom{\|x\|_1}{k}} \\
&= \frac{1}{\sigma 2^n}\sum_{x\in A}\sum_{\substack{y\in A\\ y\prec x}}\frac{1}{\binom{\|x\|_1}{\|x-y\|_1}} \\
&= \frac{1}{\sigma 2^n}\sum_{y\in A}\sum_{\substack{x\in A\\ x\succ y}}\frac{\binom{n-\|y\|_1}{\|x-y\|_1}}{\binom{\|x\|_1}{\|x-y\|_1}}\cdot\frac{1}{\binom{n-\|y\|_1}{\|x-y\|_1}} \\
&\geq \min_{\substack{x\succ y\\ x,y\in L_{\text{mid}}}}\left\{\frac{\binom{n-\|y\|_1}{\|x-y\|_1}}{\binom{\|x\|_1}{\|x-y\|_1}}\right\}\mathop{\mathbf{E}}_{\boldsymbol{y}\in A}\left[\text{score}^{\uparrow}(\boldsymbol{y},A)\right]=\Omega(\varepsilon^4)\mathop{\mathbf{E}}_{\boldsymbol{y}\in A}\left[\text{score}^{\uparrow}(\boldsymbol{y},A)\right],
\end{aligned}$$

where the final equality holds by the first part of Lemma 5.2. This proves (9), which together with

Lemma 5.4 gives

$$\begin{aligned}\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}\left[\boldsymbol{y}\in A \mid \boldsymbol{x}\in A\right] &= \frac{1}{\Theta(\sqrt{n\ln(1/\varepsilon)})}\mathop{\mathbf{E}}_{\boldsymbol{x}\in A}\left[\mathrm{score}^{\uparrow}(\boldsymbol{x},A)+\mathrm{score}^{\downarrow}(\boldsymbol{x},A)\right]\\ &= \frac{O(\varepsilon^{-4})}{\Theta(\sqrt{n\ln(1/\varepsilon)})}\left(\mathop{\mathbf{E}}_{\boldsymbol{x}\in A}\left[\mathrm{score}^{\downarrow}(\boldsymbol{x},A)\right]\right). \qquad (10)\end{aligned}$$

On the other hand, by Corollary 5.3 we have

$$\begin{aligned}\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}\left[\boldsymbol{y}\in A \mid \boldsymbol{x}\in A\right] &= \frac{\mathbf{Pr}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}[\boldsymbol{x},\boldsymbol{y}\in A]}{\mathbf{Pr}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}[\boldsymbol{x}\in A]}\\ &= \frac{\mathbf{Pr}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}'}[\boldsymbol{x},\boldsymbol{y}\in A]}{\sigma} = \Omega\left(\frac{\varepsilon^4\sigma}{\ln(1/\varepsilon)}\right). \qquad (11)\end{aligned}$$

Combining (10) with (11) and rearranging completes the proof. □

Lemma 5.5 lower bounds the average downward A-score of points $x\in A$; its conclusion may be equivalently rewritten as the following sum:

$$\sum_{x\in A}\mathrm{score}^{\downarrow}(x,A) = \Omega\left(\frac{\varepsilon^8\sigma^2\sqrt{n}2^n}{\sqrt{\ln(1/\varepsilon)}}\right). \qquad (12)$$

We may express the downward A-score $\mathrm{score}^{\downarrow}(x,A)$ of a point x as a sum over $m+1$ "buckets" of exponentially increasing size:

$$\mathrm{score}^{\downarrow}(x,A) = \left(\sum_{k\in B_0}\mathrm{dens}^{\downarrow}_k(x,A)\right)+\left(\sum_{k\in B_1}\mathrm{dens}^{\downarrow}_k(x,A)\right)+\cdots+\left(\sum_{k\in B_m}\mathrm{dens}^{\downarrow}_k(x,A)\right), \qquad (13)$$

where $B_0=\{0\}$ and $B_i=\{2^{i-1},\dots,2^i-1\}$ for each $i\in[m]$ and $m=\lceil\log(n+1)\rceil$. It will be useful for us to focus on a particular bucket $\ell\in\{0,1,\dots,m\}$ such that the overall sum of $\mathrm{score}^{\downarrow}(x,A)$ in (12) has a "large" contribution from the ℓ-th bucket. A straightforward argument, exploiting the fact that there are only logarithmically many buckets, lets us achieve this without losing too much in the sum:

Corollary 5.6. *Let $\varepsilon\geq 1/n$ and $A\subseteq L_{\mathrm{mid}}$ be a set of $\sigma 2^n$ points. There exists $\ell\leq m$ such that*

$$\sum_{x\in A}\sum_{k\in B_\ell}\mathrm{dens}^{\downarrow}_k(x,A) = \Omega\left(\frac{\varepsilon^8\sigma^2\sqrt{n}2^n}{(\log n)\sqrt{\ln(1/\varepsilon)}}\right). \qquad (14)$$

Proof. This follows from (12), (13), and the fact that there are only $m+1$ many buckets. □

Corollary 5.6 gives us a lower bound on the sum of downward A-scores of points $x\in A$ coming from a certain bucket B_ℓ. Our next corollary uses this to give a lower bound on the sum of downward A-scores of points $y\in A_u$ coming from (essentially) the same bucket B_ℓ, where A_u is an "upper vertex boundary" of A in the following sense: there exists an $|A|$-sized matching M of edges (x,y) where $x\prec y$, $x\in A$ and $y\in A_u$.

Corollary 5.7. *Let $\varepsilon \geq 1/n$ and M be a matching of $\sigma 2^n$ edges in the middle layers. Let*

$$A := \{x \in \{0,1\}^n : x \prec y \text{ and } (x,y) \in M\} \quad \text{and}$$
$$A_u := \{y \in \{0,1\}^n : y \succ x \text{ and } (x,y) \in M\}$$

be the lower and upper endpoints of edges in M, respectively. For each bucket B_i, $i \in \{0,1,\ldots,m\}$, we let $B'_i := \{j+1 : j \in B_i\}$. Then there exists an integer $\ell \leq m$ such that

$$\sum_{y \in A_u} \sum_{k \in B'_\ell} \mathrm{dens}^{\downarrow}_k(y, A) = \Omega\left(\frac{2^{\ell+n}\varepsilon^8\sigma^2}{(\log n)\sqrt{n\ln(1/\varepsilon)}}\right). \tag{15}$$

Proof. By Corollary 5.6, there exists an $\ell \leq m$ such that A satisfies (14).

Next for every edge $(x,y) \in M$ we have that

$$\mathrm{dens}^{\downarrow}_{k+1}(y, A) = \mathop{\mathbf{Pr}}_{\substack{\boldsymbol{z} \prec y \\ \|\boldsymbol{z}-y\|_1 = k+1}}[\boldsymbol{z} \in A] \geq \frac{\binom{\|x\|_1}{k}}{\binom{\|y\|_1}{k+1}} \mathop{\mathbf{Pr}}_{\substack{\boldsymbol{z} \prec x \\ \|\boldsymbol{z}-x\|_1 = k}}[\boldsymbol{z} \in A] = \frac{(k+1)\cdot \mathrm{dens}^{\downarrow}_k(x, A)}{\|x\|_1 + 1}.$$

Therefore, by (14) we have

$$\begin{aligned}
\sum_{y \in A_u} \sum_{k \in B'_\ell} \mathrm{dens}^{\downarrow}_k(y, A) &= \sum_{y \in A_u} \sum_{k \in B_\ell} \mathrm{dens}^{\downarrow}_{k+1}(y, A) \\
&\geq \sum_{x \in A_u} \sum_{k \in B_\ell} \frac{(k+1)\cdot \mathrm{dens}^{\downarrow}_k(x, A)}{\|x\|_1 + 1} \\
&= \Omega\left(\frac{\varepsilon^8\sigma^2\sqrt{n}2^n}{(\log n)\sqrt{\ln(1/\varepsilon)}} \cdot \frac{2^\ell}{n}\right).
\end{aligned}$$

This completes the proof. ∎

5.3 The weighted path tester and its analysis

Given a Boolean function f, recall that a pair (x,y) of vertices is a *violated pair with respect to f if* $x \prec y$ *and* $f(x) > f(y)$. Our algorithm `weighted-path-tester` for monotonicity testing proceeds as follows:

`weighted-path-tester`:

1. Pick a point $\boldsymbol{y}$ uniformly from L_{mid}.
2. Pick $\boldsymbol{\ell} \in \{0,1,\ldots,m = \lceil \log(n+1) \rceil\}$ uniformly.
3. Pick $\boldsymbol{k} \in B'_{\boldsymbol{\ell}}$ uniformly.
4. Pick a path $\boldsymbol{p}$ uniformly from the collection of all paths going through $0^n, \boldsymbol{y}$ and 1^n, and set $\boldsymbol{x}$ to be the (unique) point on $\boldsymbol{p}$ that has $\boldsymbol{x} \prec \boldsymbol{y}$ and $\|\boldsymbol{x} - \boldsymbol{y}\|_1 = k$.
5. Reject iff $(\boldsymbol{x}, \boldsymbol{y})$ is a violated pair.

We note that an equivalent formulation of step (4) is that $\boldsymbol{x}$ is drawn uniformly from $\{z \in \{0,1\}^n : z \prec \boldsymbol{y} \text{ and } \|\boldsymbol{y} - z\|_1 = \boldsymbol{k}\}$. Below we show that if there is a $(\sigma 2^n)$-sized matching M of violated edges of f in the middle layers of the hypercube, then the tester above succeeds in finding a violated pair with probability roughly $\Omega(\sigma^2/\sqrt{n})$.

Proposition 5.8. *Let $f : \{0,1\}^n \to \{0,1\}$ and $\varepsilon \geq 1/n$. Suppose there is a $(\sigma 2^n)$-sized matching M of violated edges of f all lying in the middle layers of the hypercube. Then* `weighted-path-tester` *succeeds (i.e. samples $\boldsymbol{x}$ and $\boldsymbol{y}$ that form a violated pair with respect to f) with probability*

$$\Omega\left(\frac{\varepsilon^8\sigma^2}{(\log^2 n)\sqrt{n\ln(1/\varepsilon)}}\right). \tag{16}$$

Proof. Let A be the 1-endpoints of edges in M, and A_u be the 0-endpoints, and note that every pair $(x, y) \in A \times A_u$ satisfying $x \prec y$ is a violated pair with respect to f. Let $\mathcal{D}^w$ denote the distribution over comparable pairs $(\boldsymbol{x}, \boldsymbol{y}) \in L_{\text{mid}} \times L_{\text{mid}}$ induced by `weighted-path-tester`. Applying Corollary 5.7, we know there that exists an $\ell^* \in \{0, 1, \ldots, m\}$ such that

$$\sum_{y \in A_u} \sum_{k \in B'_{\ell^*}} \text{dens}^{\downarrow}_k(y, A) = \Omega\left(\frac{2^{\ell^*+n}\varepsilon^8\sigma^2}{(\log n)\sqrt{n\ln(1/\varepsilon)}}\right).$$

Note that conditioning on the event of $\boldsymbol{y} = y$ and $\boldsymbol{k} = k$, the probability of $\boldsymbol{x} \in A$ is $\text{dens}^{\downarrow}_k(y, A)$. Since $\boldsymbol{y}, \boldsymbol{\ell}, \boldsymbol{k}$ are all sampled uniformly, `weighted-path-tester` succeeds with probability at least

$$\begin{aligned}
\mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}^w}\left[\boldsymbol{y} \in A_u, \boldsymbol{x} \in A\right] &= \mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}^w}[\boldsymbol{y} \in A_u] \cdot \mathop{\mathbf{Pr}}_{(\boldsymbol{x},\boldsymbol{y})\leftarrow\mathcal{D}^w}\left[\boldsymbol{x} \in A \mid \boldsymbol{y} \in A_u\right] \\
&= \frac{|A_u|}{|L_{\text{mid}}|} \cdot \frac{1}{|A_u|} \sum_{y \in A_u} \frac{1}{m+1} \sum_{\ell=0}^{m} \frac{1}{|B'_\ell|} \sum_{k \in B'_\ell} \text{dens}^{\downarrow}_k(y, A) \\
&\geq \frac{1}{(m+1)|L_{\text{mid}}||B'_{\ell^*}|} \cdot \sum_{y \in A_u} \sum_{k \in B'_{\ell^*}} \text{dens}^{\downarrow}_k(y, A) \\
&= \Omega\left(\frac{2^{\ell^*+n}\varepsilon^8\sigma^2}{(\log n)\sqrt{n\ln(1/\varepsilon)}} \cdot \frac{1}{(\log n)2^{\ell^*+n}}\right) = \Omega\left(\frac{\varepsilon^8\sigma^2}{(\log^2 n)\sqrt{n\ln(1/\varepsilon)}}\right).
\end{aligned}$$

This finishes the proof. ∎

5.4 Proof of Theorem 3

Finally we combine Proposition 5.8 with the dichotomy theorem of [CS13a] to prove Theorem 3. To state the latter, we let $v2^n$ denote the total number of violated edges in f. We also let $\sigma 2^n$ denote the size of the largest matching of violated edges in the middle layers. Then we have

Theorem 11 (Theorem 2.4 of [CS13a])**.** *For any f that is ε-far from monotone, $v \cdot \sigma = \Omega(\varepsilon^2)$.*

We now prove Theorem 3.

Proof of Theorem 3. As mentioned at the beginning of Section 5, we may assume without loss of generality that $\varepsilon \geq 1/n$ since otherwise the edge tester alone succeeds with probability $\Omega(\varepsilon/n) =$

$\Omega(\varepsilon^2)$. When $\varepsilon \geq 1/n$, our tester flips a coin, runs the edge tester with probability 1/2, and runs `weighted-path-tester` with probability 1/2. Given v and σ as defined above, the success probability of the edge tester is $\Omega(v/n)$; the success probability of `weighted-path-tester` is given in (16). It follows from Theorem 11 that the average of these two is at least

$$\Omega\left(\frac{\varepsilon^4}{n^{5/6}(\log^{2/3} n)\,(\ln(1/\varepsilon))^{1/6}}\right).$$

This finishes the proof of Theorem 3. □

Acknowledgements

We thank Eric Blais for a helpful discussion that led to an improvement of our lower bound for general hypergrid domains.

References

[ACCL07] Nir Ailon, Bernard Chazelle, Seshadhri Comandur, and Ding Liu. Estimating the distance to a monotone function. *Random Struct. Algorithms*, 31(3):371–383, 2007. 1

[BBM12] Eric Blais, Joshua Brody, and Kevin Matulef. Property testing lower bounds via communication complexity. *Computational Complexity*, 21(2):311–358, 2012. 1

[BCGSM12] Jop Briët, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Monotonicity testing and shortest-path routing on the cube. *Combinatorica*, 32(1):35–53, 2012. 1, 1, 1

[BKR04] Tugkan Batu, Ravi Kumar, and Ronitt Rubinfeld. Sublinear algorithms for testing monotone and unimodal distributions. In *ACM Symposium on Theory of Computing*, pages 381–390, 2004. 1

[BO10] Eric Blais and Ryan O'Donnell. Lower bounds for testing function isomorphism. In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity*, pages 235–246, 2010. 1.1, 2.2

[BRY13] Eric Blais, Sofya Raskhodnikova, and Grigory Yaroslavtsev. Lower bounds for testing properties of functions on hypergrid domains. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:36, 2013. 1, 1

[CS13a] Deeparnab Chakrabarty and C. Seshadhri. A $o(n)$ monotonicity tester for boolean functions over the hypercube. In *ACM Symposium on Theory of Computing*, pages 411–418, 2013. (document), 1, 1, 1.2, 3, 5.1, 5.1, 5.4, 11

[CS13b] Deeparnab Chakrabarty and C. Seshadhri. Optimal bounds for monotonicity and lipschitz testing over hypercubes and hypergrids. In *ACM Symposium on Theory of Computing*, pages 419–428, 2013. 1, 1

[CS13c] Deeparnab Chakrabarty and C. Seshadhri. An optimal lower bound for monotonicity testing over hypergrids. In *APPROX-RANDOM*, pages 425–435, 2013. 1, 1

[DGL+99] Yevgeniy Dodis, Oded Goldreich, Eric Lehman, Sofya Raskhodnikova, Dana Ron, and Alex Samorodnitsky. Improved testing algorithms for monotonocity. In *Proceedings of RANDOM*, pages 97–108, 1999. 1

[EKK+00] Funda Ergün, Sampath Kannan, S. Ravi Kumar, Ronitt Rubinfeld, and Mahesh Vishwanthan. Spot-checkers. *Journal of Computer and System Sciences*, 60:717–751, 2000. Earlier version in STOC'96. 1

[Fis04] Eldar Fischer. On the strength of comparisons in property testing. *Inf. Comput.*, 189(1):107–116, 2004. 1

[FLN+02] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 474–483, 2002. (document), 1, 8

[GGL+00] Oded Goldreich, Shafi Goldwasser, Eric Lehman, Dana Ron, and Alex Samordinsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000. 1

[GGLR98] Oded Goldreich, Shafi Goldwasser, Eric Lehman, and Dana Ron. Testing monotonicity. In *IEEE Symposium on Foundations of Computer Science*, pages 426–435, 1998. 1, 1.2

[GOWZ10] Parikshit Gopalan, Ryan O'Donnell, Yi Wu, and David Zuckerman. Fooling functions of halfspaces under product distributions. In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity*, pages 223–234, 2010. 1.1

[HK08] Shirley Halevy and Eyal Kushilevitz. Testing monotonicity over graph products. *Random Struct. Algorithms*, 33(1):44–67, 2008. 1

[MORS09] Kevin Matulef, Ryan O'Donnell, Ronitt Rubinfeld, and Rocco Servedio. Testing ±1-weight halfspaces. In *Proceedings of the 13th Annual International Workshop on Randomized Techniques in Computation*, pages 646–657, 2009. 1.1

[MORS10] Kevin Matulef, Ryan O'Donnell, Ronitt Rubinfeld, and Rocco Servedio. Testing halfspaces. *SIAM Journal on Computing*, 39(5):2004–2047, 2010. 1.1

[Mos08] Elchanan Mossel. Gaussian bounds for noise correlation of functions and tight analysis of Long Codes. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 156–165, 2008. 1.1

[O'D14] Ryan O'Donnell. *The Analysis of Boolean Functions*. Cambridge Univ. Press, 2014. Preliminary version at analysisofbooleanfunctions.org. 1.3

[RRS+12] Dana Ron, Ronitt Rubinfeld, Muli Safra, Alex Samorodnitsky, and Omri Weinstein. Approximating the influence of monotone Boolean functions in $O(\sqrt{n})$ query complexity. *ACM Transactions on Computation Theory*, 4(4):11, 2012. 1

[RS09] Ronitt Rubinfeld and Rocco A. Servedio. Testing monotone high-dimensional distributions. *Random Struct. Algorithms*, 34(1):24–44, 2009. 1

[RS13] Dana Ron and Rocco A. Servedio. Exponentially improved algorithms and lower bounds for testing signed majorities. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 1319–1336, 2013. 1.1

[VV11] Gregory Valiant and Paul Valiant. Estimating the unseen: an $n/\log(n)$-sample estimator for entropy and support size, shown optimal via new CLTs. In *ACM Symposium on Theory of Computing*, pages 685–694, 2011. 1.1, 3, 3

[Yao77] Andrew Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proceedings of the 9th Annual ACM Symposium on Theory of Computing*, pages 222–227, 1977. 1.1, 2

A Reduction from hypergrid domains $[m]^n$ when m is odd

Lemma A.1. *Let $m \in \mathbb{N}$ be odd. There exists a monotone function $h : [m]^k \to \{-1, 1\}$ such that $|\{x \in [m]^k \colon h(x) = 1\}| = |\{x \in [m]^k \colon h(x) = -1\}| + 1$, and a one-to-one mapping $\Psi : h^{-1}(-1) \to h^{-1}(1)$ such that $\Psi(x) \succ x$ for all $x \in h^{-1}(-1)$.*

Proof. The function h is defined as follows:

$$h(x_1, \ldots, x_k) = \begin{cases} 1 & \text{if } x = \lceil m/2 \rceil^k \\ \operatorname{sign}(x_i - \lceil m/2 \rceil) & \text{otherwise, where } i := \min\{i \in [k] \colon x_i \neq \lceil m/2 \rceil\}. \end{cases}$$

The monotonicity of h is straightforward to verify, as is the fact that $|\{x \in [m]^k \colon h(x) = 1\}| = |\{x \in [m]^k \colon h(x) = -1\}| + 1$. The proof is complete by noticing that the mapping

$$\Psi(x_1, \ldots, x_k) = (x_1, \ldots, x_{i-1}, -x_i, x_{i+1}, \ldots, x_k), \quad \text{where } i := \min\{i \in [k] \colon x_i \neq \lceil m/2 \rceil\}$$

is a bijection between $h^{-1}(-1)$ and $h^{-1}(1) \setminus \{\lceil m/2 \rceil^k\}$. ∎

With Lemma A.1 in hand we are ready to prove the following analogue of Proposition 4.1 for hypergrid domains $[m]^n$ when m is odd. Given the monotone function h defined in Lemma A.1, let $h' : [m]^k \to \{-1, 1, \perp\}$ be the partial function where $h'(x) = \perp$ if $x = \lceil m/2 \rceil^k$, and $h'(x) = h(x)$ otherwise (and so $|\{x \in [m]^k \colon h(x) = 1\}| = |\{x \in [m]^k \colon h(x) = -1\}| = (m^k - 1)/2$).

Proposition A.2. *For all odd $m \in \mathbb{N}$ the mapping*

$$\Phi : \{\text{all functions } f : \{-1,1\}^n \to \{-1,1\}\} \to \{\text{all functions } f : [m]^{n\lceil \log n \rceil} \to \{-1,1\}\}$$

defined by (17) below satisfies the following two properties:

1. *If $f : \{-1, 1\}^n \to \{-1, 1\}$ is monotone then $\Phi[f]$ is monotone as well.*
2. *If $f : \{-1, 1\}^n \to \{-1, 1\}$ is ε-far from monotone then $\Phi[f]$ is $\Omega(\varepsilon)$-far from monotone.*

Proof. Fix $k := \lceil \log n \rceil$. For every $f : \{-1,1\}^n \to \{-1,1\}$, we define $\Phi[f] : [m]^{kn} \to \{-1,1\}$ to be the following function: for all $x^1, \dots, x^n \in [m]^k$

$$\Phi[f](x^1, \dots, x^n) := f\big(h(x^1), \dots, h(x^n)\big). \tag{17}$$

Since h is monotone it follows that $\Phi[f]$ is monotone if f is monotone, and so it remains to show that $\Phi[f]$ is $\Omega(\varepsilon)$-far from monotone if f is ε-far from monotone. Since f is ε-far from monotone, we have by Theorem 8 that there exist $\varepsilon 2^n$ many pairwise disjoint pairs $(x^i, y^i) \in \{-1,1\}^n \times \{-1,1\}^n$ that are violations with respect to f; we will exhibit $\Omega(\varepsilon\, m^{kn})$ many pairwise disjoint pairs in $[m]^{kn}$ that are violations with respect to $\Phi[f]$, which along with another application of Theorem 8 completes the proof.

Using the same notation as in the proof of Proposition 4.1, we define the set-valued function S mapping $x \in \{-1,1\}^n$ to subsets of $[m]^{kn}$ as follows:

$$S(x) = h'^{-1}(x_1) \times \cdots \times h'^{-1}(x_n) \subseteq [m]^{kn}.$$

Note that

$$|S(x)| = \left(\frac{m^k - 1}{2}\right)^n = \frac{m^{kn}}{2^n}\left(1 - \frac{1}{m^k}\right)^n \geq \frac{m^{kn}}{2^n}\left(1 - \frac{1}{n^{\log m}}\right)^n = \Omega\left(\frac{m^{kn}}{2^n}\right)$$

for all $x \in \{-1,1\}^n$ (where we have used our choice of $k = \lceil \log n \rceil$ for the inequality), and $S(x) \cap S(y) = \emptyset$ if $x \neq y$. Furthermore, $\Phi[f](x') = f(x)$ for all $x \in \{-1,1\}^n$ and $x' \in S(x)$. In words, S maps each 1-input of f to a set of $((m^k-1)/2)^n$ many 1-inputs of $\Phi[f]$, and likewise each 0-input of f to a set of $((m^k-1)/2)^n$ many 0-inputs of $\Phi[f]$.

For any pair $(x, y) \in \{-1,1\}^n \times \{-1,1\}^n$, $x \prec y$, that is a violation with respect to f, consider pairing the $((m^k-1)/2)^n$ elements of $S(x)$ with the $((m^k-1)/2)^n$ elements of $S(y)$ via Ψ from Lemma A.1 as follows: each $a \in S(x)$, which we will view as $a = (a_1, \dots, a_n) \in ([m]^k)^n$, is paired with the unique element $b = (b_1, \dots, b_n) \in S(y)$ where $b_i = a_i$ if $x_i = y_i$, and $b_i = \Psi(a_i)$ if $x_i < y_i$. Since $x \prec y$, it follows from the definitions of S and Ψ that every $x' \in S(x)$ is paired with $y' \in S(y)$ where $x' \prec y'$. Furthermore, as noted above $\Phi[f](x') = f(x) = 1$ whereas $\Phi[f](y') = f(y) = 0$, and so every pair $(x', y') \in S(x) \times S(y)$ is a violation with respect to $\Phi[f]$. Therefore each of the $\varepsilon 2^n$ many pairs $(x, y) \in \{-1,1\}^n \times \{-1,1\}^n$ that are violations with respect to f gives rise to $\Omega(m^{kn}/2^n)$ many pairwise disjoint pairs $(x', y') \in S(x) \times S(y)$ that are violations with respect to $\Phi[f]$. Finally recalling that $S(x) \cap S(y) = \emptyset$ if $x \neq y$, we conclude that there are indeed $\varepsilon 2^n \cdot \Omega(m^{kn}/2^n) = \Omega(\varepsilon\, m^{kn})$ many pairwise disjoint pairs that are violations with respect to $\Phi[f]$. This finishes the proof. ∎

Proposition A.2 along with Theorem 1 implies the existence of a universal constant $\varepsilon_0 > 0$ such that any non-adaptive ε_0-tester for the monotonicity of $f : [m]^N \to \{-1,1\}$, where $N := n\lceil \log n \rceil$ and m is odd, must make $\tilde{\Omega}(n^{1/5}) = \tilde{\Omega}(N^{1/5})$ many queries. This along with Proposition 4.1 (establishing the same lower bound for hypergrid domains $[m]^n$ where m is even) completes the proof of Theorem 2.